

Structural Accessibility and Spectral Discreteness in Pure SU(3) Yang–Mills:

A UNNS Reassessment of Vacuum Non-Crossability and the Mass Gap

UNNS Substrate Project

October 2026

Abstract

We investigate whether the Yang–Mills mass gap requires a structurally non-crossable neighborhood around the vacuum. An initial UNNS non-crossability program distinguishes massive from gapless controls but fails to establish a universal refinement-stable vacuum boundary. We then test the alternative hypothesis that the Yang–Mills vacuum is structurally accessible within the same admissible domain. Across connected chromoelectric-field profiles, common-refinement analysis, canonical UNNS structural chambers, an unsmoothed connected-correlator control, and a vacuum-subtracted continuum energy–momentum-tensor representation, no universal positive structural exclusion floor is detected. Independent continuum glueball spectroscopy nevertheless retains a strictly positive scalar excitation scale. The combined result supports the coexistence of structural vacuum accessibility with spectral discreteness and therefore separates structural proximity to the vacuum from spectral proximity to the vacuum.

On the mathematical side, we prove a *structural–spectral separation theorem*. A positive spectral gap does not by itself imply a structural exclusion neighborhood; continuity of the excitation-energy functional at the vacuum provides a sufficient bridge to such an exclusion, while structural accessibility in the limit sense, under a persistent gap, forces that continuity to fail. The theorem is elementary—it is essentially continuity—and we say so; its role is to locate the missing assumption of the original program. We separate throughout an *empirical* accessibility statement, which the frozen evidence supports, from a *limit* statement, which it does not establish. The substantive open problem is the specialization to Yang–Mills: constructing the structure-to-spectrum map $\Pi : \mathcal{S}_{\text{UNNS}} \rightarrow \mathcal{V}_{\text{spec}}$ and determining its topology. We do not prove $V \in \overline{\mathcal{A}_{\text{YM}}}$, and nothing here addresses the Clay mass-gap problem.

Keywords: Yang–Mills mass gap; glueball spectrum; UNNS; structural accessibility; spectral gap; common refinement; falsifiability.

Contents

1	Introduction: does a mass gap require a structural gap?	3
2	Context: the mass gap, confinement, and Euclidean spectral structure	5
3	Methods and structural formalism	6
4	The original vacuum non-crossability program	10
5	From non-crossability to structural accessibility	16
6	A structural–spectral separation theorem	16

7	Stage 2: testing structural vacuum accessibility	19
8	Stage 3: independent spectral discreteness	25
9	BRIDGE-III: coexistence rather than identification	27
10	Discussion	29
11	Implications and the specialization problem	33
12	Limitations and falsifiability	34
13	Conclusion	35
A	Exhaustive verdict record	36
B	Supporting data	39
C	Reproducibility and provenance	40

1 Introduction: does a mass gap require a structural gap?

The Yang–Mills mass gap is a statement about the Hamiltonian spectrum: the lowest excited eigenvalue lies strictly above the vacuum energy, $E_1 - E_0 > 0$. The UNNS (Unbounded Nested Number Sequences) program studies recursion-generated structure, and one natural question within it is whether this spectral separation has a *structural* shadow—whether the vacuum sits behind a wall in structure space that every non-vacuum admissible configuration must respect. This paper reports the outcome of testing that idea in two successive, fully frozen research projects, and extracts from the outcome a precise mathematical statement about what such a wall would and would not follow from.

The first project, UNNS_YM_GAP, asked whether the mass gap corresponds to a refinement-stable structural non-crossability boundary. It built an explicit diagnostic, tested it against deliberately constructed gapless adversaries, and followed it through model and measured-data refinement. The idea was not established, and the closing evidence was uncertainty-limited rather than decisive. The second project, UNNS_VACUUM_CONTINUITY, then asked the converse question: is the vacuum structurally accessible, and can that accessibility coexist with an independently retained spectral scale? Its answer, within a frozen evidence architecture, is yes. Taken together the two projects are not two reports but one argument with a correction in the middle, and this paper presents them that way.

1.1 Four notions of nearness

The phrase “close to the vacuum” hides several distinct notions, and the paper is organized around keeping them apart.

- *Configuration-space nearness*: a gauge-field configuration close to the trivial one in some norm.
- *Observable-space nearness*: a gauge-invariant profile (a field strength, a correlator, an energy–momentum-tensor component) whose vacuum-subtracted contrast is small. This is what the structural chambers of the two projects actually measure; we write d_{struct} or D_{struct} for such distances.
- *Hilbert-space nearness*: a state (ray) close to the vacuum ray, e.g. in the norm of phase-fixed unit representatives.
- *Spectral nearness*: $d_{\text{spec}}(E_n, E_0) = E_n - E_0$, a property of Hamiltonian eigenvalues.

Nothing forces these to share a topology. The original hypothesis implicitly identified the second with the fourth; Section 6 shows what a sufficient bridge of this kind would give, and Section 10 returns to the four notions.

1.2 Hypotheses: isolation, empirical accessibility, limit accessibility

Let $\mathcal{A}_{\text{YM}}$ be the admissible Yang–Mills-derived structural domain, V the vacuum object, and D_{struct} a structural distance. The first project tested the *isolation* (non-crossability) hypothesis,

$$H_{\text{iso}} : \quad D_{\text{struct}}(X, V) \geq \delta > 0 \quad \text{for all admissible non-vacuum } X, \text{ stably under refinement.} \quad (1)$$

The second project replaced it by accessibility. The word “accessibility” is used in two strengths, and the distinction is central:

$H_{\text{acc}}^{\text{emp}}$: in the tested representations no universal positive exclusion floor is detected and structural vacuum accessibility is supported, while a positive spectral scale is retained; (2)

$H_{\text{acc}}^{\text{lim}}$: $\exists X_n \in \mathcal{A}_{\text{YM}} \setminus \{V\} : D_{\text{struct}}(X_n, V) \rightarrow 0$ while $E_1 - E_0 > 0$. (3)

Equation (2) is a statement about the outcome of frozen experiments; Eq. (3) is a mathematical limit statement about a sequence in a representation-independent metric. The projects support the first and do not establish the second:

H_{iso} is not established; $H_{\text{acc}}^{\text{emp}}$ is supported; $H_{\text{acc}}^{\text{lim}}$ is not yet established.

(4)

In particular the Stage-2 “structural states” are spatial profile windows and ladders of vacuum-subtracted observables, not excitation sequences carrying an energy functional (Remark 4).

1.3 What this paper establishes, proves, leaves open, and does not claim

We keep four kinds of statement strictly separate; the whole manuscript is organized around the separation. The first three are the empirical result, the elementary theorem and the open problem; the fourth records what is deliberately not claimed.

Established (empirical)	Proved (elementary theorem)	Open (UNNS problem)	Not claimed
Structural accessibility is supported ($H_{\text{acc}}^{\text{emp}}$) and coexists with two independent positive continuum scalar scales in the tested pure-SU(3) representations (Sections 7 to 9).	A spectral gap alone does not imply structural exclusion; continuity, or a local Hölder upper bound of energy by structural distance, is a sufficient bridge for such an exclusion, and accessibility plus a persistent gap rules out a continuous bridge (Section 6).	Construct the structure-to-spectrum map $\Pi : \mathcal{S}_{\text{UNNS}} \rightarrow \mathcal{V}_{\text{spec}}$ and determine its topology (Section 11).	$V \in \overline{\mathcal{A}_{\text{YM}} \setminus \{V\}}$; the limit statement $H_{\text{acc}}^{\text{lim}}$; a Clay mass-gap theorem; a derivation of the gap from UNNS structure.

1.4 Roadmap and conventions

Section 2 situates the problem; Section 3 explains the instruments (chambers, representations, decision rules) so that the empirical sections can be read without the project records; Section 4 presents the original non-crossability program in full; Section 5 explains the pivot; Section 6 proves the separation theorem; Sections 7 to 9 give the empirical body; Sections 10 and 11 interpret it; Section 12 states limitations and falsifiers; Section 13 concludes. A condensed table of every frozen verdict is given in Table 9, and the exhaustive record in Section A. Verdict strings are quoted exactly as frozen. Evidence classes follow the projects’ own convention: Class B is published measured data reconstructed from figures; Class C is exact model control anchored to a published fitted mass. Results are graded as *frozen empirical verdict*, *limited empirical verdict*, *structural interpretation*, or *not established / not claimed*.




Figure 1: Synthesis of the argument. Each box carries the frozen verdict of the corresponding stage; the arrows are logical, not causal. The last box is the output of the deterministic evaluator BRIDGE-III over frozen statuses: structural accessibility is *supported* (no universal floor detected) and a positive spectral scale is retained. It is not a statement that a limiting sequence has been constructed, and not a theorem about Yang–Mills.

2 Context: the mass gap, confinement, and Euclidean spectral structure

The mass-gap problem. The Clay Millennium formulation asks for a mathematically rigorous quantum Yang–Mills theory on $\mathbb{R}^4$ for any compact simple gauge group, together with a mass gap $\Delta > 0$ [1]. A rigorous construction in four dimensions is not available, and nothing in this paper contributes to it. The lattice formulation introduced by Wilson [9] provides a nonperturbative regularization in which the spectrum can be computed numerically [10, 11]. The projects analyzed here work entirely with published lattice and gradient-flow results for pure SU(3).

Glueball spectroscopy. In pure gauge theory the low-lying single-particle spectrum above the vacuum consists of glueballs (multi-glueball states form a continuum above them), and the lightest glueball is the scalar 0^{++} . Systematic lattice calculations of the glueball spectrum date back decades; anisotropic-lattice studies [13, 15], SU(N) comparisons [14, 12], and the high-precision continuum extrapolations of Athenodorou and Teper [16] are the anchors used here. Masses are conventionally quoted in units of the string tension, $M/\sqrt{\sigma}$, or of the Sommer scale, r_0M [18]; these are different native coordinates, and the paper never converts between them.

Mass gap is not confinement. The area law and a positive string tension σ [9, 19, 20] describe the large-distance behavior of static sources; the mass gap is the statement that the Hamiltonian has an isolated vacuum eigenvalue. The two are logically distinct statements about the same theory, even though in pure SU(3) both are controlled by a single dynamically generated scale, which is why dimensionless ratios such as $M_{0^{++}}/\sqrt{\sigma}$ are meaningful. The flux-tube profiles of [21, 22, 23] used in Stage 2 probe the gauge-field distribution around static sources; they are confinement-sector observables used here as vacuum-subtracted structural profiles, and no claim about the confinement mechanism is made.

Euclidean correlators and positivity. In a reflection-positive theory [28, 29, 30] the Euclidean two-point function of a gauge-invariant operator has the spectral representation $C(t) = \int_0^\infty \rho(E) e^{-Et} dE$ with $\rho \geq 0$ after vacuum subtraction, and the mass gap in that channel is the lower edge of $\text{supp } \rho$. Positivity of ρ makes C completely monotone; on an equally spaced

time grid it makes the correlator a Hausdorff moment sequence, which yields necessary conditions in terms of Hankel matrices [31, 32]. These facts are what the correlator-native chamber of Section 4 exploits. In practice, excited-state contamination is controlled by variational methods [27].

Gradient flow and the energy–momentum tensor. The Yang–Mills gradient flow [25] gives smoothed, finite gauge-invariant composite operators, and the energy–momentum tensor can be defined through the flow [26]. Yanagihara et al. [24] used it to measure the stress-tensor distribution around a static quark–antiquark pair; the vacuum-subtracted profile they publish is the independent representation used in Stage 2.

Model comparison. Floor-versus-zero model competition uses the Akaike criterion [33] with the small-sample correction of Hurvich and Tsai [34]; see also [35]. Interpolation sensitivity is probed with the monotone piecewise cubic Hermite interpolant (PCHIP) of Fritsch and Carlson [36].

2.1 Vacuum structure, confinement mechanisms, and spectral generation

Several programs describe the vacuum of pure Yang–Mills theory through its microscopic structure. In the monopole picture, abelian projection and dual superconductivity [2, 3] attribute confinement to condensation of magnetic monopoles; in the center-vortex picture, thick center vortices disorder Wilson loops, and lattice studies of center dominance probe this directly [4, 20]; instantons and, at finite temperature or non-trivial holonomy, calorons provide a semiclassical description [5, 6, 7]; and functional approaches study the infrared behavior of Green’s functions [8]. These programs ask *which microscopic configurations produce confinement and a dynamical scale*.

The present manuscript does not propose another mechanism of that kind and does not compare itself to them. Its question is logically prior and narrower: whether the structural nearness of a state to the vacuum, in whatever representation one adopts, must share a topology with the spectral nearness of its energy to the vacuum energy. The answer given (Section 6) is that it need not, unless a bridge such as continuity is supplied. A mechanism of any of the above kinds could in principle supply that bridge or show it to fail; this is why the structure-to-spectrum map Π of Section 11 is stated independently of any mechanism.

3 Methods and structural formalism

This section collects the machinery shared by both projects. It is deliberately self-contained: a reader unfamiliar with UNNS should be able to understand what each chamber measures, what the preregistered decision rules were, and why the instruments are complementary rather than interchangeable.

3.1 Structural objects, ladders and refinement

The basic structural object is a *ladder*: an ordered one-column list of real numbers $L = (x_1, \dots, x_n)$ extracted from a physical system—a set of glueball masses, an effective-mass series, a spatial profile sampled at increasing distance, a window of amplitudes. Two kinds of ladders appear. A *sorted* ladder is ordered by value and analyzed for its gap structure; an *ordered profile window* keeps its physical (spatial) ordering and is never sorted by amplitude. *Refinement* means passing through an ordered sequence of finer descriptions of the same physical object: finer lattice spacings at (approximately) fixed physical volume, or higher-resolution samplings of the same profile. *Common refinement* means that different refinement levels are compared only on the

intersection of their physically supported coordinate ranges, on a grid, interpolation rule, window width and minimum point count frozen in advance. A level that loses the required physical region fails admission; the grid is not moved to keep it. Extrapolation beyond measured support is prohibited.

3.2 Preregistration, freezing and evidence classes

Both projects were run falsification-first. Hypotheses, observables, decision rules, thresholds, metrics, interpolation choices, and the exact chamber builds (identified by SHA-256 digests) were frozen *before* the target data were inspected, and verdict vocabulary was fixed in advance. Negative and indeterminate outcomes are preserved under their frozen labels and are not retuned or rescued. A later record may supersede a frozen assumption only by an explicit versioned note given before the primary result is inspected. Evidence is graded: Class C denotes exact controls built from a published fitted mass (a model control, not data); Class B denotes published measured information (e.g. effective-mass trajectories) reconstructed from figures, without raw correlator matrices or full covariance. Neither class is allowed to stand in for the other; Section 4 shows why this separation proved decisive.

3.3 STRUC-I: perturbative admissibility

STRUC-I (canonical build v1.0.4) tests whether a sorted ladder remains *admissible* under perturbation of its entries. Given L , it computes the gaps $\delta_i = x_{i+1} - x_i$ and their median δ_{med} , and evaluates $K = 40$ perturbation scales κ spaced logarithmically in $[0.01, 1]$. At each scale it perturbs every entry by an independent uniform amount in $[-\varepsilon, \varepsilon]$ with $\varepsilon = \kappa \delta_{\text{med}}$, counts the order inversions produced, and compares the count with a combinatorial vulnerability capacity determined by the gaps narrower than 2ε . The reported admissibility coefficient $A_\kappa \in [0, 1]$ is the fraction of Monte Carlo draws satisfying the inequality (inversions $\leq$ capacity), averaged over scales; the structural pressure ρ is the ratio of mean inversions to capacity. Regimes are assigned by frozen boundaries: *Geometric Persistence* for mean $A_\kappa \gtrsim 0.93$, with the sub-state *Stable Structure* for the most conserved ladders, otherwise a weaker regime. If a ladder falls within 0.01 of a boundary it is rerun at 10,000 Monte Carlo draws (the setting used for the Battelli ladders below). STRUC-I is not modified anywhere in this work.

3.4 STRUC-PERC-I: percolation

STRUC-PERC-I builds a graph over the ladder entries, connecting values that lie within a κ -scaled tolerance, and reports the giant-component ratio, the number of isolated vertices and the connectivity threshold κ_{connect} . FULL_PERCOLATION requires a giant ratio of exactly 1 and no isolated vertices; otherwise the outcome is a fragmented class. It is a connectivity diagnostic and, as the first project learned (Section 4.2), connectivity is not the same thing as admissibility or as vacuum separation.

3.5 CORR-I: a correlator-native chamber

CORR-I takes a Euclidean correlator on an equally spaced grid and asks two questions: is it consistent with a positive spectral representation, and how much additional low-energy spectral weight can be inserted before the correlator is materially altered?

Baseline layer. CORR-I tests positivity of $C(t)$, monotone decrease, discrete log-convexity, the effective mass $m_{\text{eff}}(t) = \log[C(t)/C(t + \Delta t)]$, plateau persistence, and positivity of Hankel/Hausdorff moment matrices on the equally spaced grid. The verdict vocabulary is deliber-

ately limited to POSITIVE-SPECTRAL CONSISTENT, SPECTRAL-POSITIVITY TENSION, NONUNIFORM GRID --- MOMENT TEST LIMITED and INVALID FOR POSITIVE-SPECTRAL TEST.

Adversarial layer. The adversary admixes a low-energy exponential,

$$C_{\text{adv}}(t) = (1 - \alpha) C(t) + \alpha e^{-\mu(t-t_0)}, \quad (5)$$

and, rather than fixing an arbitrary α , the chamber computes the largest compatible admixture $\alpha_{\text{max}}(\mu)$ under the active deviation limits (descriptive relative limits, because pointwise covariance is unavailable). The first design used a fixed-amplitude admixture and was replaced after self-testing showed a 5% admixture already in tension at the top of the scan; the α_{max} envelope was the corrected design. An identity control $\alpha = 0$ must reproduce the baseline exactly.

Stage coordinate. The compact descriptive metric is

$$R_{\text{contract}} = \frac{\alpha_{\text{max}}(\mu/m_{\text{ref}} = 10^{-3})}{\alpha_{\text{max}}(\mu/m_{\text{ref}} = 1)}, \quad (6)$$

small values meaning strong suppression of added spectral weight as the trial scale approaches 0^+ . It is a UNNS diagnostic coordinate, not the physical gap. The margin against the hardest gapless adversary is $S = \log_{10}(R_{p=8}/R)$, called S_{fixed} when the reference is the fixed finest-resolution $p = 8$ value, S_{match} when it is the β -matched value (Class C), and S_B for the measured-data (Class B) reconstruction. $S > 0$ puts the signature on the massive side; $S \leq 0$ means the proposed boundary is reached or crossed.

3.6 VREF-I: the vacuum-referenced profile representation

VREF-I was frozen in Stage 1 as the primary representation independent of CORR-I. It does *not* use α_{max} , the Gamma family, R_{contract} , S_B , plateau fitting or Hankel acceptance; CORR-I may be compared with VREF-I but cannot define, tune or rescue it.

Vacuum proxy. For refinement level r and observable O the vacuum-referenced contrast is $\Delta O_X^{(r)}(x) = \hat{O}_X^{(r)}(x) - \hat{O}_{\text{vac}}^{(r)}(x)$; for published connected or explicitly vacuum-subtracted profiles the published convention is retained and $\Delta O = 0$ is the operational vacuum reference, $V_r^* \equiv 0$. A dataset is eligible for a primary verdict only if an explicit no-source vacuum expectation exists or the subtraction convention is documented. A published zero baseline is never treated as error-free.

Dimensionless coordinate and windows. With a reference length ℓ_{ref} declared in advance, $z(u) = \ell_{\text{ref}}^{d_O} \Delta O(x)$ with $u = x/\ell_{\text{ref}}$ and d_O the canonical mass dimension of O . Profiles are *not* normalized to their own peak, norm, integral or first point, because self-normalization would erase exactly the amplitude information needed for an approach to the vacuum. The structural object is the fixed-width ordered window $L_r(s) = (z(s + \xi_1), \dots, z(s + \xi_m))$, with the vacuum proxy $V_r^* = (0, \dots, 0)$ on the same window. Fixed width prevents the distance from shrinking artificially because fewer points remain at large radius.

Primary distance.

$$D_{\text{VREF}}^{(r)}(s) = \left[\frac{1}{m} \sum_{j=1}^m (z^{(r)}(s + \xi_j))^2 \right]^{1/2}, \quad (7)$$

the root-mean-square window distance to the vacuum proxy; an L_∞ version is allowed only as a preregistered sensitivity check. Measurement uncertainty is propagated around the distance




Figure 2: Schematic (not data) of the VREF-I construction. A vacuum-subtracted profile $z(u)$ (blue) decays toward the zero vacuum proxy. A fixed-width window of m ordered samples slides outward; $D_{\text{VREF}}(s)$ is the RMS of z over the window. The outcome of interest is whether D_{VREF} keeps decreasing toward the vacuum-compatible region or saturates at a positive floor δ_V (dashed).

but never used to redefine it, so that a noisier level cannot look closer to the vacuum merely through larger error bars.

Path and the non-triviality rule. The structural path is *spatial*: windows $L_r(s_1), L_r(s_2), \dots$ with increasing distance from the non-vacuum core toward an independently identified vacuum-like region. The following do *not* count as evidence: multiplying a fixed profile by an imposed $\lambda \rightarrow 0$; renormalizing profiles to force a vanishing tail; deleting nonzero points as the vacuum is approached; shrinking the window with s ; choosing the baseline from the same tail because it minimizes D_{VREF} ; tuning any parameter after the outcome is seen. This clause prevents the hypothesis from becoming a tautology under arbitrary amplitude attenuation.

Refinement and admissibility. A primary verdict that supports continuity (or exclusion) requires at least three admitted refinement levels on common support; a two-level sequence is only a sensitivity or replication branch and can return only a calibration-type or indeterminate verdict, never a continuity-supporting one. This is why the two-refinement CEA branch (Section 7.2) is frozen as REFINEMENT-INDETERMINATE and not promoted. D_{VREF} is not an admissibility test: a profile enters $\mathcal{A}_{\text{YM}}$ only if it independently satisfies provenance and representation rules (Yang–Mills-derived provenance, gauge-invariant observable, valid vacuum reference, known coordinate ordering and scale mapping, finite reproducible values, frozen common support, no target-dependent rescue transformation). A small D_{VREF} cannot make an invalid object admissible. STRUC-I may be run on the ordered windows as a secondary *grammar* readout, but it defines neither the vacuum proxy nor the distance.

3.7 EMT-I: the vacuum-subtracted energy–momentum-tensor representation

EMT-I changes the physical representation itself, from connected chromoelectric-field amplitudes to the gradient-flow energy–momentum tensor around a static pair. From the published profile at fixed source separation R it forms the four components $\varepsilon = -T_{44}$, $\sigma_{zz} = -T_{zz}$, $\sigma_{rr} = -T_{rr}$, $\sigma_\theta = -T_{\theta\theta}$ (sign conventions as plotted in the source, preserved alongside the derived values) and the four-component norm. The branch is analyzed on the transverse tail $r \geq R/2$.

Table 1: The instruments, what they measure, and what they cannot say.

Instrument	Measures	Cannot say
STRUC-I	Order-admissibility of a ladder under perturbation; regime	Distance to the vacuum; spectral content
STRUC-PERC-I	Connectivity of the ladder’s gap graph	Admissibility; vacuum specificity
CORR-I	Compatibility of a correlator with added low-energy weight; massive vs. gapless separation at a resolution	Hamiltonian spectrum; refinement stability by itself
VREF-I	Vacuum-referenced profile distance along a spatial path	Mass gap; asymptotic limit $D \rightarrow 0$
EMT-I	Independent vacuum-subtracted continuum representation	Same
ΔAICc	Zero floor vs. positive floor on a tail	Anything beyond the tested tail

3.8 Floor-versus-zero model competition

Whether a vacuum-directed tail decays to zero or levels off at a positive floor is decided by model competition on the tail. Two families are used, a pure exponential $Ae^{-r/\lambda}$ and a Bessel- K_0 form, each in a *zero* version ($k = 2$ parameters) and a *positive-floor* version with an additive constant δ ($k = 3$). With Gaussian errors $\text{AICc} = \chi^2 + 2k + 2k(k+1)/(n-k-1)$ for n tail points, and the decision statistic is

$$\Delta\text{AICc} = \text{AICc}_{\text{floor}} - \text{AICc}_{\text{zero}}. \quad (8)$$

Positive values favor the zero floor; the preregistered decisive threshold is $|\Delta\text{AICc}| \geq 6$; values below it are UNDECIDED. A *drop-first* sensitivity removes the innermost tail point. This criterion is conservative against a positive-floor claim as well as against a zero-floor claim: an indeterminate result is not evidence for either.

3.9 Complementary roles

Table 1 summarizes what each instrument measures. The key point is that they probe different projections of the same physical object and fail in different ways, which is why agreement across them is informative and why none is allowed to be tuned using another.

4 The original vacuum non-crossability program

The first project, UNNS_YM_GAP (closed 2026-09-30 at stage YM-3B; terminal chamber CORR-I v0.1.2), began from the frozen statement that *gapless Yang–Mills continuation is structurally non-generable under refinement, whereas the massive branch is refinement-stable*. Its stronger form was that the mass gap corresponds to a refinement-stable non-crossability boundary separating the vacuum from every non-vacuum admissible state. We present it in full, because the second project exists only because of how this one ended.

4.1 Logic of the program

Let V denote the vacuum and X an admissible non-vacuum structure. The structural hypothesis required $D_{\text{YM}}(X, V) \geq \delta_G > 0$ surviving refinement. A gapless continuation would instead be a sequence $V \rightarrow X_1 \rightarrow X_2 \rightarrow \dots$ with $E(X_n) - E(V) \rightarrow 0^+$ in which every X_n remains admissible. The intended physical bridge was

$$E(X) - E(V) \geq F_G(D_{\text{YM}}(X, V)), \quad F_G(\delta_G) > 0, \quad (9)$$

Table 2: STRUC-PERC-I outcomes over the $4\times$ deformation families (counts of FULL/HARD fragmentation outcomes, YM-2A).

Family	FULL	HARD
offset	19	61
upctrl	21	59
scale	80	0
gaponly	8	72

a structural-to-Hamiltonian bound that was never reached; the project remained a structural investigation throughout. Four non-negotiable conditions were frozen: (i) gauge-invariant observables; (ii) *adversarial* construction of the gapless continuation rather than inference from the absence of low-energy states; (iii) survival of any obstruction under common refinement; (iv) a strict distinction between a structural diagnostic and a Hamiltonian theorem. Condition (iii) was decisive: a fixed-resolution separation was never sufficient.

4.2 YM-2A: vacuum-anchored spectral ladders

The first representation took five glueball-state coordinates available at all eight Athenodorou–Teper ensembles [16]—the A_1^{++} ground state (the 0^{++} mass-gap state), the A_1^{++} first excitation, the E^{++} and T_2^{++} grounds, and the A_1^{-+} ground—and added an explicit vacuum node, $L^{\text{YM}} = (0, x_1, \dots, x_5)$. The vacuum node was essential: without it a common translation of all positive masses leaves adjacent gaps unchanged and is invisible to a gap-based chamber. Four frozen deformation families were then applied along a dyadic sequence $\epsilon = 2^{-k}$, $k = 0, \dots, 10$: *offset* (translate all positive states downward, vacuum fixed), *upctrl* (translate upward), *scale* (multiply all positive states by ϵ), and *gaponly* (move only the lowest state toward the vacuum).

STRUC-I. It did not reject the vacuum-compressed branch. Across baselines and the offset family every case remained in Geometric Persistence, with several Stable-Structure/Weak-Persistence changes but no regime transition; the six-node ladder stayed structurally admissible under a factor-1024 compression of the vacuum gap.

STRUC-PERC-I. It was more sensitive, but the directionality test failed (Table 2). Downward vacuum-gap compression often fragmented the pairwise gap graph; so did *upward* translation away from the vacuum; uniform scaling left the normalized connectivity geometry invariant; and gap-only deformation fragmented most strongly. The operative variable was selective displacement of gap coordinates relative to the internal gap distribution under the chamber’s IQR normalization, not approach to the vacuum. The correct conclusions were that spectral-gap connectivity can collapse while STRUC-I admissibility survives, and that the six-node spectral representation does not establish vacuum-specific non-crossability. This was a successful falsification of a representation/test pair, not of the program.

4.3 Why a correlator-native chamber

A short sorted list of extracted masses is already a compressed product of the Euclidean correlation structure. Euclidean two-point functions, $C(t) = \int_0^\infty \rho(E) e^{-Et} dE$ with $\rho \geq 0$, retain more direct information about spectral support. The move from STRUC-I/STRUC-PERC-I to CORR-I (Section 3.5) was therefore a correction of the representation, not a relaxation of the hypothesis.

4.4 CORR-I physical-side inputs and the legitimacy of the adversaries

AT20 does not tabulate the correlator matrix $C_{ij}(t)$; it presents effective-mass plots and the resulting mass estimates. Two controlled physical-side inputs at the finest ensemble ($\beta = 6.50$, 38^4) were therefore built, both explicitly labelled as not being the authors' raw correlators.

Figure-7 proxy. The 0^{++} ground-state effective-mass series of AT20's Fig. 7 was converted into a normalized correlator through $C(0) = 1$, $C(n+1) = C(n) e^{-aM_{\text{eff}}(n+1/2)}$. The reconstruction (11 positive points, no monotonicity violation, four log-convexity violations, no plateau, failed higher-rank Hankel tests) was classified SPECTRAL-POSITIVITY TENSION. This was not read as a failure of spectral positivity: AT20 itself discusses the late-time rise of the effective mass as a statistical fluctuation, and exact moment tests are strict enough that a noisy digitized proxy is not expected to pass them. Nevertheless its envelope contraction was strong, $R_{\text{raw}} \approx 0.00716$.

Published-mass plateau control. A clean control $C_{\text{massive}}(t) = e^{-0.3474t}$, $t/a = 0, \dots, 10$, was built from the published $aM = 0.3474(22)$, with the ± 0.0022 mass uncertainty kept as a correlated model band and not converted into independent pointwise errors. It passed every baseline test and gave $\alpha_{\text{max}}(1) = 0.5$, $\alpha_{\text{max}}(10^{-3}) \approx 0.0038438$, hence $R_{\text{massive}} \approx 0.00768765$. The agreement $R_{\text{raw}} \approx 0.00716$ versus $R_{\text{massive}} \approx 0.00769$ is the first robustness result: the late-time fluctuation affects the exact positivity tests but does not create the strong near-zero contraction.

Gapless adversaries. A single gapless control was constructed with $\rho(E) = \theta^{-1} e^{-E/\theta}$, whose spectral support reaches $E = 0$, giving $C(t) = 1/(1 + \theta t)$; the parameter $\theta = e^{0.3474} - 1$ was chosen so that the first effective-mass interval matches the massive control, $m_{\text{eff}}(0.5) = 0.3474$, so that the comparison does not simply contrast different initial scales. This control was positive, monotone, log-convex, Hankel/Hausdorff consistent and without plateau, with $R_{p=1} \approx 0.3269$, about $42.5\times$ the massive value. One control is not enough, so the family

$$C_p(t) = (1 + \theta_p t)^{-p}, \quad \theta_p = e^{0.3474/p} - 1, \quad (10)$$

was used, the Laplace transform of the Gamma density $\rho_p(E) = E^{p-1} e^{-E/\theta_p} / (\Gamma(p) \theta_p^p)$. Every member is exactly gapless (support reaching $E = 0$), is matched to $m_{\text{eff}}(0.5) = 0.3474$, and lives on the same finite interval; the family varies shape while holding the initial mass scale fixed. Three properties make this family the appropriate adversary and not a strawman. (i) Legitimacy: every member is a Laplace transform of a positive measure, so it satisfies exactly the positivity, monotonicity, log-convexity and Hankel/Hausdorff conditions that the baseline layer tests; all five members passed them. (ii) Interpolation: since $\theta_p \sim 0.3474/p$ as $p \rightarrow \infty$, $C_p(t) \rightarrow e^{-0.3474t}$, so the family converges to the massive control, and increasing p makes the adversary progressively *harder*; $p = 8$ is the hardest member tested. (iii) Falsifiability: a gapless correlator that looks exponential over a finite window can imitate massive behavior, so no single short correlator can give a binary theorem; what can be asked is whether a measurable separation persists.

4.5 YM-2 fixed-resolution result and verdict

At the finest resolution the contraction ratios were 0.500947, 0.326916, 0.178612, 0.084534 and 0.037668 for $p = 0.5, 1, 2, 4, 8$ against 0.007688 for the massive control (ratios $65.2\times$, $42.5\times$, $23.2\times$, $11.0\times$ and $4.90\times$), with the hardest gapless member still admitting about $1.60\times$ more added spectral weight than the massive control at the deepest scan point ($\alpha_{\text{max}}^{\text{deep}} = 0.006139$ versus 0.003844) (Fig. 3). All gapless members were POSITIVE-SPECTRAL CONSISTENT; hence positive spectral consistency alone does not force the massive-side response, and the near-zero contraction




Figure 3: CORR-I contraction ratio R_{contract} for the matched gapless Gamma family at $\beta = 6.50$ (blue) and for the massive plateau control (dashed). The ordering $R_{0.5} > R_1 > R_2 > R_4 > R_8 > R_{\text{massive}}$ holds; as p grows the adversary approaches the massive response but $R_{p=8} \approx 4.90 R_{\text{massive}}$ remains separated.

is not a trivial consequence of positivity. The stage closed with the provisional verdict **VACUUM-BOUNDARY DIAGNOSTIC ESTABLISHED; REFINEMENT TEST REQUIRED**. The idea thus had genuine empirical motivation: CORR-I did discriminate massive from legitimate gapless controls.

4.6 YM-3A: model-refinement persistence (Class C)

The first refinement test used exact mass-anchored exponential controls built from the eight published masses of AT20 [16], with β -matched Gamma controls at each level. The frozen result was **MODEL-REFINEMENT PERSISTENT**. At all eight levels the massive control stayed on the massive side of the fixed hardest gapless reference ($S_{\text{fixed}} > 0$) and of every β -matched control ($S_{\text{match}} > 0$); the full matched ordering was retained, all forty matched controls were positive-spectral consistent, Hankel/Hausdorff consistent and passed the $\alpha = 0$ identity null; and the native-grid window diagnostic did not reverse the separation. The matched margin grew from $S_{\text{match}} = 0.8067$ at $\beta = 5.6924$ to 1.5892 at $\beta = 6.50$ (Fig. 4). The claim boundary was explicit: an exponential built from a published fitted mass carries no measured effective-mass fluctuation, contamination, finite-sample structure or covariance. YM-3A therefore established persistence for a model family, not for measured correlators, and the project prohibited promoting it into a measured-data or continuum claim.

4.7 YM-3B: the measured-data challenge (Class B)

Source and reconstruction. The decisive test used four ordered pure-SU(3) ensembles of Sakai and Sasaki [17] at $\beta = 6.20, 6.40, 6.71, 6.93$: Wilson plaquette action, approximately fixed physical volume $La \approx 1.6$ fm, gauge-invariant spacelike Wilson-loop glueball operators, and a variational analysis with a 6×6 correlation matrix. The $A_1^{++}/0^{++}$ ground-state effective-mass trajectories (their Figs. 9–12) were digitized and reconstructed into normalized variational-eigenvalue proxies on a common physical-time grid $u = 0.31, 0.36, \dots, 0.61$ (seven points). All four levels passed input admission. Interpolation (linear primary, PCHIP sensitivity), matched Gamma adversaries, nested ladders and the uncertainty procedure were frozen before the scientific batch.

Central result. The central margins were $S_B = +0.020350, +0.035867, -0.003395, +0.051851$. At $\beta = 6.71$ the matched ordering failed because $R_B > R_{p=8}$: the central trajectory *crossed* the




Figure 4: YM-3A (Class C): fixed-reference margin S_{fixed} and β -matched margin S_{match} against β for the eight AT20 ensembles. Both are positive at every level (the two coincide at $\beta = 6.50$ by construction).

Table 3: Class B central margins S_B with the preregistered 95% intervals. Only $\beta = 6.71$ is sign-unstable.

β	S_B (central)	95% interval	sign-stable
6.2	0.020350	[0.004053, 0.035659]	yes
6.4	0.035867	[0.017167, 0.053094]	yes
6.71	-0.003395	[-0.027454, 0.019320]	no
6.93	0.051851	[0.024726, 0.081194]	yes

proposed boundary (Table 3; full detail in Table 12, Section B).

Robustness of the crossing. The negative sign survived the PCHIP interpolation (-0.003948) and all three nested ladder lengths ($S_B^{L5} = -0.007663$, $S_B^{L6} = -0.006666$, $S_B^{L7} = -0.003395$; Fig. 6). It therefore could not be dismissed as an artifact of the interpolation rule or of the choice of five, six or seven common-grid points.

Why it is not a falsification. The preregistered 2000-draw pointwise-no-covariance ensemble at $\beta = 6.71$ gave a 95% interval $[-0.027454, +0.019320]$, median -0.003285 , with 60.65% of draws negative and 39.35% positive, while the other three levels were sign-stable (Fig. 5). Under the frozen rule digitization/statistical uncertainty dominates the sign of the only central crossing and the verdict is DATA-INDETERMINATE. A further limitation: all four reconstructed baselines were classified SPECTRAL-POSITIVITY TENSION under the strict moment checks, consistent with their provenance as digitized trajectories rather than covariance-preserving correlators, and the $\beta = 6.93$ reconstruction did not satisfy the plateau detector even though its margin remained positive. The crossing is not erased; it is not promoted beyond the precision of the source.

4.8 Closure and what it taught

The central hypothesis was recorded as NOT ESTABLISHED. YM-4 (a bridge theorem under the original premise) was not entered, and further tests (Liang et al. replication, extension to Chen et al., additional source search, CORR-I redesign to restore the separation) were cancelled. The record does *not* call the picture refuted, does not claim that no structural boundary exists, and neither proves nor disproves the physical mass gap. It left five lessons that carry directly into




Figure 5: Class B margin $S_B = \log_{10}(R_{p=8}/R_B)$ at the four Sakai–Sasaki ensembles [17]. Dots: central values; bars: preregistered 2000-draw 95% intervals (2.5%–97.5% quantiles); dashed line: the proposed boundary $S_B = 0$. Only $\beta = 6.71$ is sign-unstable, and there the interval straddles zero.




Figure 6: Sensitivity of the central margin to the nested common-grid ladder length (L5, L6, L7 points) and to PCHIP interpolation. At $\beta = 6.71$ the sign is negative in every variant.

the second project: connectivity is not admissibility; representation depth matters, since a short sorted mass ladder erases distinctions that the Euclidean correlator retains; positive spectral consistency does not determine the near-zero response; a finite-window gapless spectrum can approach massive-looking behavior continuously; and model-level persistence does not imply measured-data persistence, so evidence classes must stay separated.

It also left one consequence at the level of UNNS ontology. The picture implicit in the original program was

$$\text{admissible structure} \mid \text{forbidden region} \mid V,$$

and the experiments did not support that geometry. The more faithful reading is

$$\text{admissible structure} \longrightarrow \text{increasingly vacuum-like structure} \longrightarrow V$$

inside a common structural domain: the vacuum behaves as an internal structural regime or locus of the substrate and not as an exterior beyond it. This is a structural interpretation, not a theorem about quantum field theory, and it is the starting point of the second project.

5 From non-crossability to structural accessibility

The pivot is conceptual. A mass gap separates *eigenvalues*; a structural neighborhood describes *configurations or observables*. These live in different spaces with no a priori common topology, so a spectral gap need not induce a structural wall. The second project, UNNS_VACUUM_CONTINUITY, froze the corresponding ontology in its Stage 1: the vacuum is an *internal closure object* of the admissible domain,

$$V \in \overline{\mathcal{A}_{\text{YM}} \setminus \{V\}}, \quad \text{primary hypothesis} \quad \inf_{X \in \mathcal{A}_{\text{YM}} \setminus \{V\}} D_{\text{struct}}(X, V) = 0, \quad (11)$$

against the isolation alternative: $\exists \delta_V > 0$ with $D_{\text{struct}}(X, V) \geq \delta_V$ for every sufficiently refined admissible non-vacuum structure in the tested representation class. Stage 1 deliberately did *not* adopt the language “basin”, “attractor” or “fixed point”, which would require further evidence. Its physical shorthand, frozen as interpretive and not theorem-level, is that the mass gap separates eigenvalues, not necessarily structural neighborhoods.

Frozen requirements. Within VREF-I, evidence for structural vacuum continuity required all of: at least three admitted refinement levels; a documented independent vacuum baseline or explicit connected subtraction; common physical support; fixed-width ordered windows; decreasing or vacuum-compatible D_{VREF} toward the independently identified asymptotic region; no refinement-stable positive lower bound resolved above uncertainty; persistence under the preregistered metric and interpolation sensitivities; no loss of admissibility; and no dependence on CORR-I. A finite dataset can support structural continuity *over its tested domain*; it cannot by itself prove the infinite-resolution statement $D_{\text{struct}} \rightarrow 0$.

Frozen falsification condition. Within the tested representation class, the continuity hypothesis is contradicted if sufficiently refined admitted data establish a stable nonzero exclusion scale $D_{\text{VREF}} \geq \delta_V > 0$ through the full vacuum-directed region, surviving common refinement, uncertainty propagation, the frozen secondary metric and independent profile construction. If the vacuum reference is inadequate, the common grid fails, or uncertainties dominate the putative bound, the verdict is *indeterminate*, not isolation-supported. Seven specific failure modes F1–F7 were enumerated in advance (Table 6).

Separation from the spectral side. Three quantities are kept distinct: D_{VREF} , the structural profile distance; Δ_{spec} , a measured or reconstructed positive spectral scale; and Δ_{YM} , the physical continuum Yang–Mills mass gap. Stage 2 addresses the first; Stage 3 tests whether the first coexists with the second in the same theory.

6 A structural–spectral separation theorem

Section 5 suggests that the original program was implicitly assuming a bridge from spectral energy to structural distance. This section proves what such a bridge would give, and what its absence allows.

6.1 Setting

Definition 1 (Structural–spectral system). A *structural–spectral system* $(\mathcal{S}, d, V, \mathcal{P}, \epsilon)$ consists of a metric space $(\mathcal{S}, d)$ of structural states; a vacuum $V \in \mathcal{S}$; a set $\mathcal{P} \subseteq \mathcal{S} \setminus \{V\}$ of *physical*

non-vacuum excitations; and an excitation-energy functional $\varepsilon : \mathcal{S} \rightarrow [0, \infty)$ with $\varepsilon(V) = 0$. It is gapped with gap $\Delta > 0$ if

$$\varepsilon(X) \geq \Delta \quad \text{for all } X \in \mathcal{P}. \quad (12)$$

It has a *structural exclusion radius* $\delta > 0$ if $d(X, V) \geq \delta$ for all $X \in \mathcal{P}$, i.e. $B_\delta(V) \cap \mathcal{P} = \emptyset$ for the open ball $B_\delta(V) = \{X : d(X, V) < \delta\}$. It is *structurally accessible* if there is a sequence $X_n \in \mathcal{P}$ with $d(X_n, V) \rightarrow 0$.

Thus H_{iso} asserts an exclusion radius, and $H_{\text{acc}}^{\text{lim}}$ asserts structural accessibility, both under a gap. The set $\mathcal{S}$ may contain non-physical states; this matters in Remark 1.

Theorem 1 (Structural–spectral separation). *Let $(\mathcal{S}, d, V, \mathcal{P}, \varepsilon)$ be gapped with gap $\Delta > 0$.*

- (a) **A gap alone does not imply a structural gap.** *There exist gapped systems with $\inf_{X \in \mathcal{P}} d(X, V) = 0$; so $\Delta > 0$ does not imply an exclusion radius.*
- (b) **Continuity is a sufficient bridge.** *If ε is continuous at V , then there is $\delta > 0$ with $B_\delta(V) \cap \mathcal{P} = \emptyset$.*
- (c) **Accessibility forces discontinuity.** *If there are $X_n \in \mathcal{P}$ with $d(X_n, V) \rightarrow 0$, then ε is discontinuous at V ; its oscillation there satisfies $\limsup_{X \rightarrow V} |\varepsilon(X) - \varepsilon(V)| \geq \Delta$.*
- (d) **Factorization obstruction.** *Let $(\mathcal{Y}, \tau)$ be a topological space (a “spectral-state space”), $\Pi : \mathcal{S} \rightarrow \mathcal{Y}$ and $\mathcal{E} : \mathcal{Y} \rightarrow [0, \infty)$ with $\varepsilon = \mathcal{E} \circ \Pi$, and put $\Omega = \Pi(V)$, so $\mathcal{E}(\Omega) = 0$. If $\mathcal{P}$ contains a sequence $X_n \rightarrow V$ in $(\mathcal{S}, d)$, then at least one of the following fails: Π is continuous at V (into τ); $\mathcal{E}$ is continuous at Ω (in τ).*
- (e) **Quantitative form.** *Suppose there are $r > 0, C > 0, \alpha > 0$ with $\varepsilon(X) \leq C d(X, V)^\alpha$ whenever $d(X, V) < r$. Then every $X \in \mathcal{P}$ satisfies*

$$d(X, V) \geq \delta_{\text{struct}} := \min \left\{ r, \left(\frac{\Delta}{C} \right)^{1/\alpha} \right\}. \quad (13)$$

If the Hölder bound holds globally ($r = \infty$), $\delta_{\text{struct}} = (\Delta/C)^{1/\alpha}$.

Proof. (a) Take $\mathcal{S} = \{0\} \cup \{1/n : n \in \mathbb{N}\} \subset \mathbb{R}$ with the Euclidean metric, $V = 0$, $\mathcal{P} = \{1/n : n \in \mathbb{N}\}$, and $\varepsilon(0) = 0, \varepsilon(1/n) = 1$ for all n . Then $\varepsilon \geq \Delta := 1$ on $\mathcal{P}$, so the system is gapped; yet $d(1/n, V) = 1/n \rightarrow 0$, so $\inf_{X \in \mathcal{P}} d(X, V) = 0$ and no exclusion radius exists.

(b) Put $\eta = \Delta/2$. By continuity at V there is $\delta > 0$ with $d(X, V) < \delta \Rightarrow |\varepsilon(X) - \varepsilon(V)| < \Delta/2$, i.e. $\varepsilon(X) < \Delta/2$ because $\varepsilon(V) = 0$. If some $X \in \mathcal{P}$ had $d(X, V) < \delta$, then $\Delta \leq \varepsilon(X) < \Delta/2$, impossible since $\Delta > 0$. Hence $B_\delta(V) \cap \mathcal{P} = \emptyset$.

(c) Suppose ε were continuous at V . Then $X_n \rightarrow V$ would give $\varepsilon(X_n) \rightarrow \varepsilon(V) = 0$, contradicting $\varepsilon(X_n) \geq \Delta > 0$. For the oscillation: $|\varepsilon(X_n) - \varepsilon(V)| = \varepsilon(X_n) \geq \Delta$ along a sequence converging to V .

(d) Suppose, for contradiction, that Π is continuous at V and $\mathcal{E}$ is continuous at $\Omega = \Pi(V)$. A composition of maps continuous at the relevant points is continuous at V , so $\varepsilon = \mathcal{E} \circ \Pi$ is continuous at V in $(\mathcal{S}, d)$, contradicting (c). (Only the definition of continuity at a point is used; no countability assumption on τ is needed.)

(e) Let $X \in \mathcal{P}$ and suppose $d(X, V) < \delta_{\text{struct}}$. Then $d(X, V) < r$, so $\varepsilon(X) \leq C d(X, V)^\alpha$; and $d(X, V) < (\Delta/C)^{1/\alpha}$ gives $C d(X, V)^\alpha < \Delta$. Hence $\varepsilon(X) < \Delta$, contradicting (12). So $d(X, V) \geq \delta_{\text{struct}}$. $\square$

Remark 1 (Continuity is sufficient, not necessary). Part (b) does not say that exclusion *requires* continuity of ε . Take $\mathcal{S} = \mathbb{R}$, $V = 0$, $\mathcal{P} = [1, 2]$, $\varepsilon = 1$ on $\mathcal{P}$, $\varepsilon = 5$ on $(0, \frac{1}{2})$ and $\varepsilon = 0$ elsewhere. The system is gapped ($\Delta = 1$) and has exclusion radius 1, yet ε is discontinuous at V , because non-physical states in $(0, \frac{1}{2})$ carry energy. Continuity at V therefore provides a *sufficient* bridge

from a gap to exclusion; the accurate summary is: a positive gap does not by itself imply a structural exclusion neighborhood; continuity of ε at V supplies such an implication; and accessibility under a persistent gap rules the continuity out.

Remark 2 (Sharpness of (13)). On $\mathcal{S} = [0, \infty)$, $V = 0$, $\varepsilon(x) = Cx^\alpha$ and $\mathcal{P} = \{x \geq (\Delta/C)^{1/\alpha}\}$ the system is gapped with gap Δ and its exclusion radius equals $(\Delta/C)^{1/\alpha}$ exactly, so the constant cannot be improved in general.

Corollary 1 (Topological reading). *Let $d_\varepsilon(X, Y) = |\varepsilon(X) - \varepsilon(Y)|$ be the energy pseudometric. If the system is gapped and structurally accessible, the identity map $(\mathcal{S}, d) \rightarrow (\mathcal{S}, d_\varepsilon)$ is discontinuous at V : the set $\{\varepsilon < \Delta/2\}$, a d_ε -neighborhood of V , is not a d -neighborhood of V . The structural topology therefore fails to be finer than the energy topology at the vacuum.*

Proof. $\{\varepsilon < \Delta/2\}$ contains V and is open for d_ε but contains no point of $\mathcal{P}$, whereas every d -ball about V contains some $X_n \in \mathcal{P}$ for large n . $\square$

6.2 Illustrations

Example 1 (Amplitude-vanishing gapped correlators). Let $\mathcal{S}_\Delta$ consist of the zero correlator $V \equiv 0$ together with all $C_\nu(t) = \int_{[\Delta, \infty)} e^{-Et} d\nu(E)$, $t \geq 0$, ν a nonzero finite positive measure, with $d(C, C') = \sup_{t \geq 0} |C(t) - C'(t)|$. Set $\varepsilon(C_\nu) = \inf \text{supp } \nu \geq \Delta$, $\varepsilon(V) = 0$, $\mathcal{P} = \mathcal{S}_\Delta \setminus \{V\}$. Since C_ν is decreasing, $d(C_\nu, V) = C_\nu(0) = \nu([\Delta, \infty))$. For $\nu_n = \frac{1}{n} \delta_\Delta$ we get $d(C_{\nu_n}, V) = 1/n \rightarrow 0$ while $\varepsilon(C_{\nu_n}) = \Delta$ for all n . This system is gapped and structurally accessible, and ε is discontinuous at V by Theorem 1(c).

Example 2 (A Hilbert-space instance). Let H be a self-adjoint operator on a Hilbert space $\mathcal{H}$ with a simple eigenvalue $E_0 = 0$ (eigenvector Ω) and $\sigma(H) \setminus \{0\} \subset [\Delta, \infty)$. For a unit vector ψ write $\psi_\perp = \psi - \langle \Omega, \psi \rangle \Omega$ and let $\mu_\perp$ be the spectral measure of H in $\psi_\perp$; define $\varepsilon(\psi) = \inf \text{supp } \mu_\perp$ if $\psi_\perp \neq 0$ and $\varepsilon(\psi) = 0$ otherwise, and let $\mathcal{P}$ be the set of unit vectors with $\psi_\perp \neq 0$. Then $\varepsilon \geq \Delta$ on $\mathcal{P}$. For $\psi_\theta = \cos \theta \Omega + \sin \theta \varphi$ with $\varphi \perp \Omega$, $\|\varphi\| = 1$, in the spectral subspace of $[\Delta, \infty)$, we have $\psi_\theta \in \mathcal{P}$, $\|\psi_\theta - \Omega\| \rightarrow 0$ as $\theta \rightarrow 0$, and $\varepsilon(\psi_\theta) \geq \Delta$ for every θ ; moreover $|\langle \psi_\theta, O \psi_\theta \rangle - \langle \Omega, O \Omega \rangle| \rightarrow 0$ for every bounded operator O . Here accessibility holds in the norm and observable topologies while the support-edge functional ε is norm-discontinuous at the vacuum. This is a generic illustration of the theorem and says nothing about how UNNS structure maps into a Yang–Mills Hilbert space.

The first example is a toy, not a Yang–Mills statement. It shows that the answer to “does a gap imply a wall?” depends on the structural metric: an amplitude-sensitive distance sees spectral weight tending to zero as the vacuum is approached, whereas a shape-normalized distance of the kind a ratio-based diagnostic such as R_{contract} uses would not assign the vacuum a point in the space at all. The second shows that the same phenomenon occurs in the Hilbert-space topologies in which one would hope to formulate the Yang–Mills problem. Choosing the topology on $\mathcal{S}_{\text{UNNS}}$ is part of the open problem.

6.3 Status: what is and is not claimed

Remark 3 (Do not oversell). Theorem 1 is essentially a statement about continuity at a point; it is mathematically valid and complete as stated but elementary, and we do not present it as a major theorem. Its value is diagnostic: it makes the missing premise of the original program explicit. H_{iso} was effectively the implication $\Delta > 0 \Rightarrow D_{\text{struct}}(X, V) \geq \delta > 0$; by (a) that implication is not automatic, and by (b) and (e) it holds under an additional bridge condition such as (13).

Table 4: Stage 2 branches at a glance.

Branch	Data	Instrument	Frozen verdict
7.1	Baker et al. [22], 3 resolutions	VREF-I/A calibration	CALIBRATION-SENSITIVE
7.2	Cea et al. [21], 2 refinements	VREF-I/B	REFINEMENT-INDETERMINATE
7.3	Cea et al., 5 ladders	STRUC-I, STRUC-PERC-I	ARCHITECTURE-PRESERVED
7.4	Battelli–Bonati [23], 3 ladders	STRUC-I, STRUC-PERC-I	UNSMOOTHED-ARCHITECTURE-CONSISTENT
7.5	Yanagihara et al. [24], $R = 0.69$ fm	EMT-I	EMT-UNCERTAINTY-LIMITED
7.6	Same, $R = 0.46, 0.92$ fm	EMT-I	GEOMETRY-ROBUST-NO-FLOOR

Remark 4 (The premise is not empirically instantiated). Parts (c)–(d) assume a sequence of *physical excitations* $X_n \rightarrow V$ in a representation-independent structural metric. The experiments do not establish this. The Stage-2 structural states are spatial profile windows and ladders of vacuum-subtracted observables (Section 7), not excitation sequences carrying a prescribed energy functional, and no representation-independent D_{struct} is constructed. The experiments support accessibility empirically ($H_{\text{acc}}^{\text{emp}}$); they do not discharge the premise of the theorem ($H_{\text{acc}}^{\text{lim}}$). The coexistence claim of Section 9 therefore lives inside the frozen evidence architecture and is not an application of Theorem 1 to Yang–Mills.

Remark 5 (On the space $\mathcal{Y}$). In (d) the space $\mathcal{Y}$ is deliberately abstract: “energy” is not a scalar function naturally defined on arbitrary Hilbert-space vectors, so at this stage we take $\mathcal{Y}$ to be a spectral-state/data space on which a functional $\mathcal{E}$ is specified, and reserve an actual Hilbert-space realization (such as the support-edge functional of Example 2) for the future Yang–Mills construction. We write $\mathcal{Y}_{\text{spec}}$ for this space when it is instantiated for UNNS.

7 Stage 2: testing structural vacuum accessibility

Stage 2 of UNNS_VACUUM_CONTINUITY is the main empirical body. It asks whether any admitted representation yields a refinement-stable positive structural exclusion floor around the vacuum. Success was not defined as one negative diagnostic value. The test required a combination of common refinement, vacuum-directed physical profiles, asymptotic zero-floor versus positive-floor competition, uncertainty control, independent structural chambers, smoothing sensitivity, an independent operator/representation change, geometry sensitivity, and preservation of negative and indeterminate results without retuning. The important point is not that every branch was positive; it is that the negative and limited branches repeatedly failed to produce the universal positive floor that non-crossability demands, and that those that could not decide were preserved as such.

Each branch below is described with the same structure: *data* $\rightarrow$ *construction* $\rightarrow$ *decision rule* $\rightarrow$ *result* $\rightarrow$ *sensitivity* $\rightarrow$ *what the verdict permits and forbids*. The six branches and their frozen verdicts are summarized in Table 4.

7.1 Baker 2022: VREF-I/A calibration

Data and construction. Three pure-SU(3) resolutions of the flux-tube data of Baker et al. [22]: R1 ($\beta = 6.240$, 33 rows), R2 ($\beta = 6.544$, 25 rows) and R3 ($\beta = 6.769$, 33 rows), each

mirror-symmetric in transverse distance, with native supports $[-1.022, 1.022]$ fm for R1 and $[-0.511, 0.511]$ fm for R2 and R3. All three files passed the frozen SHA-256, finite-value, three-column, non-negative-error, mirror-symmetry, ordering and common-support checks. The frozen geometry: reference length $d = 0.511$ fm, $u = x_i/d$, common grid $u = 0, 0.125, \dots, 1$, window width $W = 0.5$, window size $m = 5$, linear interpolation as primary and PCHIP as sensitivity, extrapolation prohibited.

Provenance and erratum. The inputs are the three transverse profile files of the longitudinal nonperturbative field at $x_l = d/4 \simeq 0.128$ fm (EX_NP_x10.128_beta6.240/6.544/6.769.dat) and the lattice geometry ($a(\beta)$ and $d = 0.511$ fm). They were verified byte-for-byte against the ancillary data files that the publisher identifies as the data tables of the published plots, and their SHA-256 digests match the frozen configuration (Section C). The 2023 erratum [22] corrects only the significant digits of the systematic error on $\sqrt{F}$ at $\beta = 6.769$ in the last line of the force table; it does not alter the profile data or geometry used here.

Decision rule. For each window the RMS profile discrepancies R_{12}, R_{23}, R_{13} between refinement levels are computed, and the contraction ratio $C = R_{23}/R_{12}$ asks whether the medium-fine discrepancy is smaller than the coarse-medium one. CALIBRATION-SENSITIVE is assigned if the PCHIP sensitivity changes at least one qualitative property: a profile sign, the direction signature of the D_{RMS} path, the class of C relative to 1, or the ordering of R_{12}, R_{23}, R_{13} in any window; otherwise, after integrity checks, CALIBRATION-VALID.

Result. Only window 3 ($u = 0.375$) changes qualitative class: $C = 0.987 (< 1)$ with linear interpolation and $C = 1.036 (> 1)$ with PCHIP, with the refinement-distance ordering changing as well (Fig. 7). The inner windows 0–2 have $C < 1$ under both interpolations (medium-fine discrepancy smaller than coarse-medium, interpolation-robustly), and the outermost window 4 has $C > 1$ under both (1.485 linear, 1.529 PCHIP), so the loss of refinement contraction at the largest transverse distance is not produced by the interpolator. The fixed-window distance D_{VREF} decreases monotonically outward at every level under both interpolations (R1: $1.821 \rightarrow 0.614$, R2: $1.862 \rightarrow 0.653$, R3: $1.880 \rightarrow 0.684$, i.e. -66% , -65% , -64%), but terminal values remain substantially positive.

Sensitivity diagnosis. The sensitivity is localized, not global, and it is a grid-commensurability effect in R2: the common grid spacing 0.0639 fm is aligned with every R1 native step and every second R3 step ($a \approx 0.0639, 0.0426, 0.0320$ fm for R1–R3), but many common-grid points fall midway between R2 native points. The maximum dimensionless interpolation difference is 0.0178 for R2 against 1.1×10^{-4} for R1 and 7.6×10^{-5} for R3; for R2 it is about 1.18 times the published pointwise error scale. The window 3 flip occurs because C lies almost exactly at unity. A descriptive leave-one-out diagnostic shows that interpolation curvature can exceed the pointwise error on the coarser profiles (median $|r|/\epsilon$ of 3.04, 2.17, 0.50 for R1–R3). A post hoc native-aligned diagnostic at $u \in \{0, 0.25, 0.5, 0.75, 1\}$ gives $C \approx 0.927$; it is diagnostic only and cannot alter the frozen outcome.

What the verdict permits and forbids. CALIBRATION-SENSITIVE forbids reading the branch as CONTINUITY-SUPPORTED or ISOLATION-SUPPORTED; the common support ends at 0.511 fm, so the branch does not test the asymptotic vacuum-compatible region. Its contribution is methodological: a structural-vacuum conclusion cannot rest on an interpolation-sensitive common-refinement construction. Three design lessons were frozen for the next branch: prefer natively common




Figure 7: VREF-I/A calibration: contraction ratio $C = R_{23}/R_{12}$ by window, linear versus PCHIP interpolation. The classification relative to $C = 1$ (dashed) flips only in window 3, where both values lie within 4% of unity.

coordinates, make interpolation explicit in the uncertainty if it is unavoidable, and require access farther into the vacuum-compatible tail.

7.2 Cea et al. 2017: VREF-I/B

Data and construction. The connected chromoelectric-field profile of Cea et al. [21] supplies two physical refinements (R1, R2) with a deep vacuum-directed tail and an explicit zero-field vacuum reference. The extraction is deterministic (22 rows; 9 common-grid points; common support 0 to 0.762 fm; primary tail cut $x \geq 0.380$ fm giving 5 tail points) and was hashed before analysis.

Decision rules. Three questions: is the path vacuum-directed (linear and PCHIP); does the zero-floor or positive-floor model win on the tail (Section 3.8); and do the two refinements agree on the absolute amplitude (the frozen refinement-consistency rule).

Result. The path is VACUUM-DIRECTED under both linear and PCHIP interpolation, with no interpolation sensitivity. The floor-model comparison gives $\Delta\text{AICc} = 5.238$ for both the exponential and K_0 families on the primary tail (UNDECIDED, just below 6), and 7.200 for both under the drop-first sensitivity (ZERO, crossing the decisive threshold in favor of the zero floor) (Fig. 9). The refinement check fails: $\chi_{\text{ref}}^2/N_{\text{tail}} = 46.05$ with a *structured* mismatch, the second refinement's tail amplitude exceeding the first's by factors 1.07–1.18 over $x = 0.38$ –0.76 fm (Fig. 8).

What the verdict permits and forbids. The zero-floor model is preferred qualitatively and, under drop-first, decisively; but the full preregistered branch cannot be promoted, because the amplitude mismatch means the two refinements do not agree on an absolute vacuum-distance calibration. REFINEMENT-INDETERMINATE was preserved unchanged. It is not evidence for isolation; it is evidence that the primary CEA observable does not yet provide a refinement-stable absolute calibration. The branch reports a vacuum-directed structural path with no detected positive floor.

7.3 CEA structural bridge: STRUC-I and STRUC-PERC-I

Construction. The ordered amplitude ladders for five CEA constructions (R1 native, R2 native, R1 common support, R2 common-linear, R2 common-PCHIP) were passed through the canonical




Figure 8: CEA tail amplitudes for the two physical refinements (R1, R2 common-linear) at the five common-grid tail points, and their ratio. The ratio departs systematically from unity (1.07–1.18), the structured mismatch that blocks promotion beyond REFINEMENT-INDETERMINATE.

chambers without modification.

Result and sensitivity. All five ladders returned the same regime, *Geometric Persistence / Stable Structure*, and FULL_PERCOLATION. The question was whether the absolute-amplitude mismatch propagates into a regime or topology change; it does not. The frozen synthesis is ARCHITECTURE-PRESERVED.

Meaning. Refinement changes absolute amplitude more strongly than it changes structural organization. The interpretation is deliberately narrow: it speaks about the ordered amplitude-ladder architecture, not about distance to the vacuum.

7.4 Battelli and Bonati 2019: an unsmoothed control

Purpose. To test whether the preserved architecture depends on the smearing used in the CEA construction.

Data and construction. The unsmoothed connected-correlator measurement of the longitudinal chromoelectric component in $SU(3)$ [23] (stochastically exact evaluation, no smoothing algorithm). Three independent positive-transverse-distance ladders B1, B2, B3 with $n = 4, 6, 5$ rungs; the $d_t = 0$ point was removed before ingestion under the frozen renormalization rule. STRUC-I v1.0.4 ran with 10,000 Monte Carlo draws and 40 perturbation scales in $[0.01, 1]$; STRUC-PERC-I v2.5.0 with its standard settings.

Result. All three ladders are in *Geometric Persistence / Stable Structure* (mean $A_\kappa = 0.999962, 0.999965, 0.999992$; minimum $0.9986, 0.9988, 0.9997$; maximum structural pressure $0.346, 0.373, 0.287$) and all three return FULL_PERCOLATION (giant ratio 1, no isolated vertices, $\kappa_{\text{connect}} = 2.0, 1.0, 2.0$, the values above 1 arising from the chamber’s adaptive extension and retained as returned).

What the verdict permits and forbids. UNSMOOTHED-ARCHITECTURE-CONSISTENT: there is no refinement-induced regime split, and the architecture is not restricted to the CEA smoothing. Because the ladders have only 4–6 values, the percolation outcomes remain very-short-ladder diagnostics. The branch is corroborative; it is not an independent primary vacuum-distance measurement and does not turn REFINEMENT-INDETERMINATE into CONTINUITY-SUPPORTED. (The

exported metadata fields `dataset.family` and `dataset.class` were left null; they are export labels, not chamber inputs, and have no numerical effect.)

7.5 Yanagihara 2019: an independent EMT representation

Purpose. This branch is stronger than the Battelli control: it changes the physical representation itself, from connected field amplitudes to the vacuum-subtracted continuum energy–momentum tensor.

Data and extraction. The gradient-flow EMT profile around a static quark–antiquark pair [24], panel (b), $R = 0.69$ fm, recovered *deterministically from the vector geometry of the published PDF*: marker and error-bar coordinates were read directly, with no raster digitization or OCR. Each of the four components has 185 points on *exactly the same* native transverse coordinates (740 long-format rows; support 3.3×10^{-8} to 0.658 fm), so no interpolation is needed and the PCHIP sensitivity is not applicable. The marker/error-bar midpoint discrepancy is at most 2.6×10^{-4} PDF points. The primary tail $r \geq R/2 = 0.345$ fm contains 127 native coordinates.

Decision rules. Vacuum-directed or not for the four-component norm; zero-floor versus positive-floor model competition with both families and the drop-first sensitivity; component-wise classification into NONZERO-FLOOR, UNCERTAINTY-LIMITED or NONMONOTONE; decisive threshold $|\Delta\text{AICc}| \geq 6$.

Result. The four-component norm is VACUUM-DIRECTED. The four comparisons give $\Delta\text{AICc} = +1.939$ (exponential, primary), $+2.020$ (K_0 , primary), $+1.961$ and $+2.034$ (drop-first): all favor the zero floor numerically, none reaches the threshold. Component checks: ε , σ_{rr} , σ_θ are UNCERTAINTY-LIMITED, σ_{zz} is NONMONOTONE, and no component returns NONZERO-FLOOR.

What the verdict permits and forbids. The branch remains EMT-UNCERTAINTY-LIMITED and is not promoted to EMT-CONTINUITY-CONSISTENT. It forbids neither a small positive floor nor decisive confirmation of a zero floor; it permits the statement that, in an independent representation, the vacuum-directed behavior survives and no positive floor is returned.

7.6 Geometry sensitivity

The same EMT-I analysis was run, without retuning, at $R = 0.46$ fm (tail $r \geq 0.23$ fm, 156 coordinates) and $R = 0.92$ fm (tail $r \geq 0.46$ fm, 88 coordinates). In both the four-component norm is VACUUM-DIRECTED, no geometry returns EMT-POSITIVE-FLOOR, and no component is NONZERO-FLOOR (Table 5). The twelve ΔAICc values lie in 0.82–2.09, all below 6 (Fig. 9). Because the maximum is 2.09, descriptive decisive thresholds of 4, 6 or 10 would all leave every EMT comparison undecided; this observation does not alter the frozen threshold of 6 or any frozen verdict. The synthesis GEOMETRY-ROBUST-NO-FLOOR is

$$R = 0.46, 0.69, 0.92 \text{ fm} \implies$$

vacuum-directed EMT structure with no detected positive structural floor,

with the explicit qualification that all three geometries remain uncertainty-limited with respect to a decisive asymptotic zero-floor proof. It does not establish that the Hamiltonian mass gap vanishes, that $D_{\text{struct}} \rightarrow 0$ is a theorem, or that every Yang–Mills representation must behave this way.

Table 5: EMT-I component classifications by source separation (four-component norm: VACUUM-DIRECTED at all three). UL = UNCERTAINTY-LIMITED, NM = NONMONOTONE.

R (fm)	$\epsilon = -T_{44}$	σ_{zz}	σ_{rr}	σ_θ
0.46	UL	UL	UL	UL
0.69	UL	NM	UL	UL
0.92	NM	UL	UL	UL




Figure 9: Zero-floor versus positive-floor model competition ($\Delta\text{AICc} > 0$ favors the zero floor). Dashed line: the frozen decisive threshold $|\Delta\text{AICc}| = 6$. Only the CEA drop-first sensitivity crosses it; all EMT-I entries (three source separations, primary and drop-first tails, two families) lie below it, so the data there neither exclude nor establish a positive floor.

7.7 Synthesis and falsification audit

The Stage 2 synthesis is NO-UNIVERSAL-STRUCTURAL-EXCLUSION-DETECTED, with the structural interpretation STRUCTURAL-VACUUM-ACCESSIBILITY-SUPPORTED. Concretely: (1) vacuum-directed behavior survives the move from connected field profiles to a vacuum-subtracted continuum EMT; (2) the EMT result survives the change of source separation without a floor; (3) ordered structural architecture is preserved across the CEA refinements even where absolute amplitude agreement fails; (4) the same architecture appears in the unsmoothed Battelli control; (5) no frozen primary or secondary branch returns a robust ISOLATION-SUPPORTED, EMT-POSITIVE-FLOOR, or equivalent positive-floor result. Table 6 records the outcome of the seven preregistered failure modes.

What Stage 2 has not established. It does not prove $V \in \overline{\mathcal{A}_{\text{YM}}}$. It has not delivered, with decisive precision, a single fully refinement-stable measured-data sequence $L_n \in \mathcal{A}_{\text{YM}}$ with $D_{\text{struct}}(L_n, V) \rightarrow 0$ under a representation-independent metric: the strongest primary CEA branch is REFINEMENT-INDETERMINATE and the primary EMT branch EMT-UNCERTAINTY-LIMITED. The absence of a detected floor is, in those branches, absence of decisive evidence for or against

Table 6: Stage 2 falsification audit (F1–F7), as recorded in the Stage 2 synthesis.

	Failure criterion	Outcome
F1	Every admissible refinement path terminates before vacuum approach	Not observed (CEA and EMT contain vacuum-directed paths over their measured domains)
F2	Common refinement produces $D_{\text{struct}} \geq \delta > 0$	Not observed (no stable positive-floor verdict)
F3	Apparent approach disappears under independent representations	Not observed
F4	Approach requires loss of admissibility	Not observed in the structural chambers used; not proved for all possible representations
F5	Effect produced only by digitization, interpolation or window choice	Not supported (Baker downgraded; EMT needs no interpolation)
F6	Continuity-like behavior only in synthetic controls	Not observed (measured lattice and published continuum profiles)
F7	An independent structural chamber reproduces universal non-crossability	Not observed

one; the support for accessibility comes from the convergence of branches and representations, not from any single decisive test. This is exactly the distinction between $H_{\text{acc}}^{\text{emp}}$ (supported) and $H_{\text{acc}}^{\text{lim}}$ (not established) of Eq. (4).

8 Stage 3: independent spectral discreteness

Stage 3 asks a different question of independent data: is there a strictly positive spectral scale for the lowest scalar excitation of pure SU(3) gauge theory? The stage was deliberately designed so that the structural and spectral sides stay independent: the structural side is carried forward unchanged from Stage 2, and the spectral side is established independently from continuum glueball spectroscopy. The target is not a numerical relation between them but coexistence. Two continuum anchors were frozen, each in its *native coordinate*, and deliberately not converted into one another.

8.1 SPEC-I/A: Athenodorou–Teper

The continuum 0^{++} ground state in units of the string tension, taken from [16], is

$$\frac{M_{0^{++}}}{\sqrt{\sigma}} = 3.405(21), \quad (14)$$

with the alternate source-supplied extrapolation 3.391(23) and lower one-sigma bounds 3.384 and 3.368; the frozen status is SPECTRAL-SCALE-POSITIVE. The underlying ladder of eight ensembles ($a\sqrt{\sigma} = 0.3999 \rightarrow 0.1038$) has $M/\sqrt{\sigma}$ rising from 2.468(42) to 3.346(23), positive at every level (Fig. 10, left; Table 15). These are lattice values, not continuum-extrapolated, shown only to display the approach to the continuum anchor. For context, Table 7 lists the other channels from the same source; only the A_1^{++} ground state enters the verdict.




Figure 10: Spectral positivity in native coordinates. *Left*: Athenodorou–Teper lattice $M_{0++}/\sqrt{\sigma}$ against $a\sqrt{\sigma}$ (error bars as recorded), with the continuum anchor 3.405(21) as a dashed line. *Right*: the Chen et al. anchor $r_0 M_{0++} = 4.16$ with its combined statistical-plus-anisotropy bar ($0.11 + 0.04$ added linearly). The panels use different coordinates and are not to be compared on one axis.

Table 7: Continuum glueball masses in units of $\sqrt{\sigma}$ extracted from [16] (Table 14 of the source); only the first row enters SPEC-I/A.

Channel	State	$M/\sqrt{\sigma}$
$A1^{++}$	ground	3.405(21)
$A1^{++}$	first excitation	5.855(41)
E^{++}	ground	4.904(20)
$T2^{++}$	ground	4.884(19)
$A1^{-+}$	ground	5.276(45)

8.2 SPEC-I/B: Chen et al.

In the independent anisotropic-lattice study [15] (five coupling values) the scalar glueball mass is

$$r_0 M_{0++} = 4.16(11)(4) \quad (15)$$

(statistical and anisotropy errors). The preregistered conservative bound subtracts the two uncertainties linearly, $4.16 - 0.11 - 0.04 = 4.01 > 0$; the frozen status is SPECTRAL-SCALE-POSITIVE. A historical source value 4.21(11) was recorded as a check and *not* counted as an independent Stage 3 branch.

8.3 Independence and the refusal to convert

The two anchors use different actions, scale conventions ($\sqrt{\sigma}$ versus the Sommer parameter r_0 [18]) and collaborations. We do not convert between $\sqrt{\sigma}$ and r_0 merely to display agreement: independence is part of the argument, and each scale is positive in its own coordinate (Fig. 10, right). Forcing the two onto a common axis would add an assumption (a value of $r_0\sqrt{\sigma}$) that neither source needs.

Table 8: Margin of the frozen spectral scales above zero, each in its own native coordinate. Lower bounds are those recorded in the SPEC-I extracts (one-sigma for A; conservative linear for B).

Branch	Coordinate	Central value	Lower bound	Central value / σ
SPEC-I/A	$M_{0++} / \sqrt{\sigma}$	3.405	3.384	$3.405/0.021 \approx 162$
SPEC-I/B	$r_0 M_{0++}$	4.16	4.01	$4.16/0.117 \approx 36$

9 BRIDGE-III: coexistence rather than identification

BRIDGE-III is a deterministic evaluator over *frozen statuses only*: it reads the Stage 2 status and the two SPEC-I statuses, verifies six checks, and emits an outcome; it does not rerun any chamber, refit any profile, alter any spectral number, or modify Stage 2. Its inputs and result are

STRUCTURAL-VACUUM-ACCESSIBILITY-SUPPORTED (Stage 2)
 $\wedge$ SPECTRAL-SCALE-POSITIVE (SPEC-I/A) $\wedge$ SPECTRAL-SCALE-POSITIVE (SPEC-I/B)
 $\implies$ COEXISTENCE-SUPPORTED.

All six deterministic checks passed (structural status found; SPEC-I/A and SPEC-I/B statuses found; both numeric guards; theory match to pure SU(3)), and the three input files matched their recorded SHA-256 digests (Section C). The record states that earlier verdicts were not modified and that neither a mass-gap theorem nor a closure theorem is claimed. The project-level frozen result is STRUCTURAL-SPECTRAL-COEXISTENCE-SUPPORTED: within the frozen pure-SU(3) evidence architecture, structural accessibility is supported and coexists with spectral discreteness. Because the structural input is $H_{\text{acc}}^{\text{emp}}$, the coexistence claim is likewise empirical; it is not the limit statement $H_{\text{acc}}^{\text{lim}}$.

(The last column divides the central value by the recorded one-sigma uncertainty, with $\sigma_B = \sqrt{0.11^2 + 0.04^2} \approx 0.117$ used for provenance only.) Both scales are many standard deviations from zero; the point of Stage 3, however, is not the size of the margin but that positivity is retained independently of the structural stage.

9.1 Master table of frozen verdicts

Table 9 condenses every frozen verdict and stage status of both projects with what each rules in and out; the exhaustive record with the representation tested and the question asked is in Section A.

Verdict glossary. Each frozen string is a decision-rule output, not a free-text assessment; the exhaustive record is Section A.

CALIBRATION-SENSITIVE a qualitative property of a common-refinement comparison changes under the interpolation sensitivity.

REFINEMENT-INDETERMINATE two physical refinements disagree structurally in absolute amplitude, so refinement stability can be neither supported nor refuted.

ARCHITECTURE-PRESERVED the structural regime and percolation topology are unchanged across refinements.

UNSMOOTHED-ARCHITECTURE-CONSISTENT the same architecture appears on an unsmoothed correlator.

EMT-UNCERTAINTY-LIMITED vacuum-directed path, but zero-floor versus positive-floor is not decisive under the frozen threshold.

GEOMETRY-ROBUST-NO-FLOOR no floor at any tested source separation.

NO-UNIVERSAL-STRUCTURAL-EXCLUSION-DETECTED no refinement-stable positive floor is detected across the tested representations.

STRUCTURAL-VACUUM-ACCESSIBILITY-SUPPORTED the empirical accessibility statement H_{acc}^{emp} .

SPECTRAL-SCALE-POSITIVE the continuum scalar scale is strictly positive in its native coordinate.

COEXISTENCE-SUPPORTED accessibility and spectral positivity coexist (BRIDGE-III).

MODEL-REFINEMENT PERSISTENT separation persists across refinement on exact model controls.

DATA-INDETERMINATE measured-data persistence is neither supported nor refuted.

NOT ESTABLISHED the hypothesis is neither proved nor refuted.

Table 9: Master table of frozen verdicts (condensed from Table 10). G = UNNS_YM_GAP, V = UNNS_VACUUM_CONTINUITY.

ID	Source	Verdict	Rules in / out
G1	UNNS_YM_GAP YM-0	complete for the present program	In: preregistered protocol. Out: any physical claim.
G2	UNNS_YM_GAP YM-1	sufficient for current work	In: data base for YM-2/3. Out: nothing physical.
G3	UNNS_YM_GAP YM-2A / STRUC-I	vacuum-compressed branch not rejected	Out: short ladder as non-crossability evidence.
G4	UNNS_YM_GAP YM-2A / STRUC-PERC-I	directionality test failed; representation/test pair falsified	Out: connectivity collapse $\equiv$ vacuum non-crossability.
G5	UNNS_YM_GAP YM-2B	VACUUM-BOUNDARY DIAGNOSTIC ESTABLISHED; REFINEMENT TEST REQUIRED (provisional)	In: fixed-resolution separation. Out: refinement stability.
G6	UNNS_YM_GAP YM-3A	MODEL-REFINEMENT PERSISTENT	In: model-family persistence. Out: measured-correlator persistence.
G7	UNNS_YM_GAP YM-3B	DATA-INDETERMINATE	In: central crossing at $\beta = 6.71$. Out: stable falsification; stable persistence.
G8	UNNS_YM_GAP YM-4	not entered / cancelled	Out: any structural-to-Hamiltonian bridge.
G9	UNNS_YM_GAP closure	NOT ESTABLISHED	Out: universal non-crossability as established; not a proof of its negation.
V0	UNNS_VACUUM_CONTINUITY Stage 1	FROZEN (definitions; no empirical verdict)	Out: "basin", "attractor", "fixed point".
V1	Baker 2022 [22] VREF-I/A	CALIBRATION-SENSITIVE	In: methodological warning. Out: continuity or isolation verdict.
V2	CEA 2017 [21] VREF-I/B	REFINEMENT-INDETERMINATE	In: vacuum-directed path, no floor detected. Out: absolute refinement-stable calibration; isolation.
V3	CEA 2017 STRUC bridge	ARCHITECTURE-PRESERVED	In: structural organization stable. Out: absolute-amplitude agreement.
V4	Battelli 2019 [23] control	UNSMOOTHED-ARCHITECTURE-CONSISTENT	In: corroboration. Out: independent primary distance measurement.
V5	Yanagihara 2019 [24] EMT-I, $R = 0.69$ fm	EMT-UNCERTAINTY-LIMITED	In: vacuum-directed, zero floor favored numerically. Out: decisive floor exclusion; EMT-CONTINUITY-CONSISTENT.
V6	Yanagihara 2019 geometry	GEOMETRY-ROBUST-NO-FLOOR	In: no floor at any tested separation. Out: floor-free proof.
V7	Stage 2 synthesis	NO-UNIVERSAL-STRUCTURAL-EXCLUSION-DETECTED	Out: universal non-crossability in tested representations.
V8	Stage 2 interpretation	STRUCTURAL-VACUUM-ACCESSIBILITY-SUPPORTED	In: accessibility (empirical). Out: $V \in \overline{\mathcal{A}_{YM}}$ as a theorem.




Figure 11: The four nearness notions as a chain of spaces and maps (schematic). The structure-to-spectrum map Π is the open object of this paper; the energy functional $\mathcal{E}_{\text{sp}}$ is continuous or not depending on the topology chosen on $\mathcal{Y}_{\text{spec}}$; the spectral gap acts only on the final energy axis. Stage 2 lives in the second box, Stage 3 on the last.

ID	Source	Verdict	Rules in / out
V9	SPEC-I/A: AT20 [16]	SPECTRAL-SCALE-POSITIVE	In: $d_{\text{spec}} > 0$ in $\sqrt{\sigma}$ units. Out: any structural statement.
V10	SPEC-I/B: Chen et al. [15]	SPECTRAL-SCALE-POSITIVE	In: independent positivity. Out: conversion between $\sqrt{\sigma}$ and r_0 .
V11	BRIDGE-III	COEXISTENCE-SUPPORTED	In: coexistence. Out: identification; mass-gap or closure theorem.
V12	Stage 3 / project	STRUCTURAL-SPECTRAL-COEXISTENCE-SUPPORTED	In: structural continuity can coexist with spectral discreteness (within frozen architecture). Out: Clay problem.

10 Discussion

10.1 Four notions of nearness

The original program asked a single question—is the vacuum surrounded by a structurally empty neighborhood?—but the word *near* was carrying four different meanings. Separating them is, we think, the main conceptual gain of the whole exercise, and it is what turns the theorem of Section 6 from an appended remark into the organizing statement of the paper.

Configuration-space nearness. Two field configurations are near if some metric on the configuration (or ensemble) space is small. This is the notion implicit in the phrase “approaching the vacuum” for a gauge field, and it is the one least accessible to our instruments: none of the frozen chambers operates on configurations. No representation-independent metric on $\mathcal{A}_{\text{YM}}$ was constructed here.

Observable-space nearness. Two states are near if a chosen family of gauge-invariant observables takes nearby values. All Stage 2 evidence lives here: the vacuum-subtracted chromoelectric profile, the energy-density tail, the four-component EMT norm. The statement “no positive floor detected” is a statement about this notion in five specific representations (Section 7.7). Observable-space nearness is representation-dependent by construction, which is why the structural chambers were required to agree across representations rather than within one.

Hilbert-space nearness. Two states are near if their rays in the physical Hilbert space are close, e.g. in the norm of representative unit vectors with the overall phase fixed (equivalently, in the Fubini–Study distance). The vacuum is the unique ground state, and states $\psi_\theta = \cos \theta \Omega + \sin \theta \varphi$ with $\varphi \perp \Omega$ converge to Ω in norm as $\theta \rightarrow 0$. This notion is complete and canonical, but it is a property of the quantized theory and is not measured by any lattice profile directly. Example 2




Figure 12: Schematic of the separation (not data). *Left*: a structural space in which a sequence X_n converges to V ($D_{\text{struct}} \rightarrow 0$). *Right*: the energy axis, with the vacuum level E_0 and a gap Δ . The map $\mathcal{E} = E - E_0$ along the sequence jumps from 0 at V to a value $\geq \Delta$ at every X_n : the sequence is structurally convergent but spectrally separated. This is the situation of Theorem 1(a),(c); nothing here asserts that Yang–Mills realizes it.

shows what is at stake: along such a norm-convergent sequence a support-edge functional can stay at a fixed positive value, so Hilbert-space nearness does not by itself force energy nearness either.

Spectral nearness. Two states are near if their energies are close, i.e. $d_{\text{spec}}(E_n, E_0) = |E_n - E_0| \rightarrow 0$. A positive gap $\Delta = E_1 - E_0$ is the statement that spectral nearness to E_0 is impossible for any state orthogonal to the vacuum. This is the only notion in which a gap has an unconditional meaning.

What the theorem says about the four. Convergence in the first three notions does not by itself imply convergence in the fourth. Continuity at the vacuum of the map from the first three to energies is a sufficient condition guaranteeing such an implication. Theorem 1(b) states that such continuity is a *sufficient* bridge to an exclusion neighborhood; Theorem 1(c) states that accessibility under a persistent gap forces the oscillation of $\mathcal{E}$ at V to be at least Δ ; and Theorem 1(d) localizes the failure to one of two maps. In plain words, the original program tacitly assumed that being close to the vacuum in structure means being close in energy; the gap is a statement about the second and says nothing about the first.

The sentence that must not be over-read. Theorem 1 establishes that structural convergence need not imply spectral convergence. The empirical results are compatible with this architecture and support structural accessibility, but do not establish the limiting sequence required to instantiate the theorem in Yang–Mills. In the notation of Eqs. (2) and (3): $H_{\text{acc}}^{\text{emp}}$ is supported; $H_{\text{acc}}^{\text{lim}}$ is not established; the premise of the theorem is therefore not empirically instantiated (Remark 4), and the coexistence claim of BRIDGE-III is empirical coexistence within a frozen architecture, not a demonstration that the abstract counterexample is realized by the theory.

10.2 Three readings of the coexistence

Suppose future work upgrades $H_{\text{acc}}^{\text{emp}}$ to $H_{\text{acc}}^{\text{lim}}$ while a positive gap remains. By Theorem 1(d), with the factorization $\mathcal{E} = \mathcal{E}_{\text{sp}} \circ \Pi$ through $\Pi : \mathcal{S}_{\text{UNNS}} \rightarrow \mathcal{Y}_{\text{spec}}$, at least one of two maps fails to be continuous at the vacuum. Three readings are logically available, and they are not equivalent.

- (R1) Π is **discontinuous** at V . Structural neighbors of the vacuum are spectrally far away: the structure-to-spectrum map jumps. This is the literal content of the abstract counterexample and the reading most directly suggested by the phrase “structurally accessible but spectrally

separated”. It predicts that a faithful Π built from gauge-invariant correlators would show a finite jump in the extracted spectral data as the structural distance goes to zero, at fixed physical scale.

- (R2) $\mathcal{E}_{\text{sp}}$ **is discontinuous at the vacuum point of $\mathcal{Y}_{\text{spec}}$** . The map from spectral states to energies is the discontinuous one. This happens if the topology of $\mathcal{Y}_{\text{spec}}$ is too weak (for instance, if it identifies states whose spectral measures differ only in their lowest support point). Then the discontinuity is a feature of how spectral states are topologized, not of the structure-to-spectrum map, and a different choice of $\mathcal{Y}_{\text{spec}}$ could remove it. The Hilbert-space example (Example 2) is of this kind: the functional $\psi \mapsto (\text{edge of the spectral support})$ is not norm-continuous at Ω along $\psi \perp \Omega$.
- (R3) **The topology on the structural side is the issue**. What we call D_{struct} in the five representations may simply be too coarse: two states with small observable distance may be far in a finer structural topology that would also see the gap. In that case “accessibility” would be an artifact of the resolution of the representation, and a sufficiently fine structural topology would restore an exclusion neighborhood. The frozen design cannot exclude this reading; the representation-diversity argument of Stage 2 reduces but does not eliminate it.

None of the three is selected by the present data. The point of the list is that they ask for *different* work: (R1) asks for an explicit construction and computation of Π ; (R2) asks for a choice of $\mathcal{Y}_{\text{spec}}$ and a proof of continuity or discontinuity of the energy functional on it; (R3) asks for a structural metric that is demonstrably not coarser than the spectral one. We regard (R2) as the most tractable, because it is a question about functional analysis on spaces of spectral measures that can be formulated without lattice input, and (R3) as the most dangerous, because it would reduce the Stage 2 conclusion to a statement about resolution.

10.3 Confinement, scale generation and the mass gap

The three terms in the title of this subsection are often run together and the frozen evidence touches them unequally.

Confinement. The flux-tube profiles of Stage 2 are confinement-sector observables. We use only their *vacuum-subtracted tail*, i.e. the statement that the field returns toward the vacuum value away from the flux tube, and we do not model the tube itself. Nothing here constrains the confinement mechanism.

Scale generation. The spectral anchors are quoted in units of $\sqrt{\sigma}$ and r_0 , two scales generated dynamically by the theory. That $M_{0++}/\sqrt{\sigma} \approx 3.4$ and $r_0 M_{0++} \approx 4.2$ are $O(1)$ numbers in these units is the standard lattice statement that the theory has a single scale; the present paper uses it only to say that a positive spectral scale is present in two independent normalizations, and does not convert between them.

Mass gap. The Clay formulation asks for a rigorous construction of the quantum theory on $\mathbb{R}^4$ together with a positive lower bound on the spectrum of the Hamiltonian above the vacuum [1]. The lattice glueball masses used here are strong numerical evidence for a positive scalar excitation scale in a lattice regularization continued to the continuum [12, 13, 15, 14, 16]; they are not a proof of the continuum gap, and they are used here only as an *independent positive input* in Stage 3. The structural stage never touches the Hamiltonian. This is the sense in which the conclusion is “coexistence” and not “derivation of the gap”.

10.4 What each stage contributes and what it cannot

Reading the whole program backwards from the theorem is instructive. Stage 1 (the original program) contributes a falsified test and a surviving diagnostic: the CORR-I statistic separates massive from gapless controls at fixed resolution, and the refinement test is indeterminate on measured data. Stage 2 contributes a robust negative: no universal positive floor was detected across profile, tail, bridge, control and geometry variations. Stage 3 contributes a positive that is independent of the structure. The theorem contributes a conditional logical remark: if the empirical accessibility result can eventually be upgraded to $H_{\text{acc}}^{\text{lim}}$, then Theorem 1 implies that the relevant structural–spectral bridge cannot remain continuous at the vacuum. Because only $H_{\text{acc}}^{\text{emp}}$ is established, the theorem does not yet force any such discontinuity; it names the object (II) that must be studied to decide.

It is worth stating what *would* be a different, stronger paper: one that (i) constructs II on a concrete class of states, (ii) proves $H_{\text{acc}}^{\text{lim}}$ in that class, and (iii) shows the energy functional is discontinuous. We do none of these. The paper established a consistent logical frame, a body of frozen empirical evidence organized by it, and a precisely stated open problem.

10.5 Methodological lessons

Several lessons from the two projects bear on how such programs should be run, independent of the Yang–Mills application.

- (M1) **Freeze before the scientific batch.** Every verdict string in Table 9 was defined with its decision rule before the data were examined, and no verdict was modified afterwards. The outcome that most deserves this protection is the uncomfortable one: $\beta = 6.71$ in YM-3B crosses zero, and the preregistered label (DATA-INDETERMINATE) is what the frame forced rather than what a narrative would have preferred.
- (M2) **Distinguish a falsified test from a falsified hypothesis.** STRUC-PERC-I falsified the *representation/test pair* (the directionality test failed); it did not falsify non-crossability. Treating the first as the second would have been an over-reach in the opposite direction.
- (M3) **Report model-family persistence and measured-data persistence separately.** A clean persistence on exactly-constructed controls (Class C) says nothing about whether the same separation survives noise, reconstruction error and finite statistics (Class B). The gap between the two classes is where the original program stalled.
- (M4) **Prefer cross-representation agreement to within-representation precision.** The positive result of Stage 2 does not come from any single fit; it comes from the same qualitative outcome in five different observables, none of which was constructed to produce it.
- (M5) **Keep numerical preference and decisive exclusion apart.** Zero-floor models are numerically favored in the EMT sector ($\Delta\text{AICc} \approx 1.9\text{--}2.1$ at the primary geometry), but the frozen decisive threshold is 6; the verdict is correspondingly EMT-UNCERTAINTY-LIMITED. Only the CEA drop-first sensitivity (7.200) crosses the threshold, and a sensitivity analysis is not allowed to promote the primary verdict.
- (M6) **Separate admissibility from distance.** A profile can approach the vacuum proxy and still fail admission; VREF-I keeps the two checks apart so that the structural distance cannot be inflated by inadmissible inputs.

11 Implications and the specialization problem

11.1 The Π problem

The specialization of the theorem to Yang–Mills is the construction of the structure-to-spectrum map $\Pi : \mathcal{S}_{\text{UNNS}} \rightarrow \mathcal{Y}_{\text{spec}}$, where $\mathcal{Y}_{\text{spec}}$ is a topological space of spectral states. A minimal requirement list follows from Section 6.

- (P1) **A concrete $\mathcal{S}_{\text{UNNS}}$.** A class of structural states closed enough that D_{struct} is a genuine (pseudo)metric rather than a chamber-specific score, and in which the vacuum V is a point whose neighborhoods can be discussed.
- (P2) **A concrete $\mathcal{Y}_{\text{spec}}$.** A space of spectral data—for instance positive spectral measures ν on $[0, \infty)$ representing Euclidean correlators $C(t) = \int e^{-Et} d\nu(E)$ [28, 29, 30]—equipped with a stated topology (weak-*, total variation, or a Hausdorff-moment topology [31, 32]). The choice decides the regularity of the energy functional $\mathcal{E}_{\text{sp}}(\nu) = \inf(\text{supp } \nu \setminus \{E_0\}) - E_0$. This support-edge functional is in general not continuous in the weak topology: arbitrarily small low-energy mass can disappear in a weak limit, so the edge can jump upward. Under the usual assumptions it is upper rather than lower semicontinuous, and its precise regularity depends on the chosen spectral-state space. This is exactly the delicacy of (R2).
- (P3) **A definition of Π on gauge-invariant correlators.** The correlator chamber CORR-I is a first concrete step: it maps a Euclidean correlator to a positivity/contraction diagnostic. It is not yet a map into $\mathcal{Y}_{\text{spec}}$.
- (P4) **A statement of what is held fixed.** Physical volume, resolution, and scale-setting ([18, 27]) must be specified in the definition of the limit; the present work only compares at matched physical scale.

Beyond (P1)–(P4), three concrete problems turn the open item into a research agenda.

- (P5) **The CORR-I envelope as a coordinate.** Study the added-weight envelope $\alpha_{\text{max}}(\mu)$ of Section 3.5 as a candidate coordinate on a space of spectral measures, and determine its continuity under refinement and under the topologies of (P2).
- (P6) **Paired structural–spectral data.** Seek datasets in which a structural representation and a spectral extraction come from the *same ensembles*, and construct a finite-dimensional candidate Π_N from them. The Athenodorou–Teper spectral ladder must not serve as the structural side, since that would make the construction circular.
- (P7) **Regularity of the support-edge functional.** Study $\mathcal{E}_{\text{sp}}$ under the weak-*, total-variation, Wasserstein and stronger constrained topologies, and determine which additional lower-mass conditions near the spectral edge (for instance a uniform lower bound on the spectral weight just above the edge) are needed for continuity.

11.2 A testable quantitative form

Part (e) of Theorem 1 gives a quantitative bridge: if $\mathcal{E}$ is α -Hölder with constant C near the vacuum, structural distance $\delta = \min\{r, (\Delta/C)^{1/\alpha}\}$ is excluded. For the program this is more useful than an abstract statement, because it turns the question into a number to be measured. If a candidate family of structural states X_n were produced with $D_{\text{struct}}(X_n, V) \rightarrow 0$ and a measured $\mathcal{E}(X_n)$, the statistic

$$\hat{C}_\alpha = \sup_n \frac{\mathcal{E}(X_n)}{D_{\text{struct}}(X_n, V)^\alpha}$$

would be a sequence-level diagnostic. A bounded $\hat{C}_\alpha$ is a consistency condition for a Hölder bridge along that sequence, but it does not establish the neighborhood-wide bound $\mathcal{E}(X) \leq C D_{\text{struct}}(X, V)^\alpha$ required by Theorem 1(e); an exclusion radius requires a uniform bound over all relevant admissible structural states sufficiently near V . Divergence rules out that Hölder exponent along the sequence. In the genuine $H_{\text{acc}}^{\text{lim}}$ -plus-gap case, where $D_{\text{struct}}(X_n, V) \rightarrow 0$ and $\mathcal{E}(X_n) \geq \Delta > 0$, divergence is automatic for every $\alpha > 0$:

$$\frac{\mathcal{E}(X_n)}{D_{\text{struct}}(X_n, V)^\alpha} \geq \frac{\Delta}{D_{\text{struct}}(X_n, V)^\alpha} \rightarrow \infty.$$

The estimator is therefore informative mainly as a falsifier of a proposed bridge, and its use requires exactly the sequence $H_{\text{acc}}^{\text{lim}}$ that the present evidence does not provide.

Illustration with the SPEC-I/A scale. Take $\Delta_*/\sqrt{\sigma} = 3.405$ (Section 8) purely as a numerical illustration of the units involved. If future work produced states X_n with $D_{\text{struct}}(X_n, V) \rightarrow 0$ and $\mathcal{E}(X_n)/\sqrt{\sigma} \geq 3.405$, then for every $\alpha > 0$ the ratio $\mathcal{E}(X_n)/\sqrt{\sigma} / D_{\text{struct}}(X_n, V)^\alpha \rightarrow \infty$, so no Hölder bridge of any exponent could hold. Conversely, a uniform local bound $\mathcal{E}(X)/\sqrt{\sigma} \leq C D_{\text{struct}}(X, V)^\alpha$ would imply $D_{\text{struct}}(X, V) \geq (3.405/C)^{1/\alpha}$ for every admissible state with that energy, which is an explicit exclusion radius. This does not identify the Stage 2 profile windows with the SPEC-I/A excitations: the two stages remain independent, and no energy is assigned to any Stage 2 structural state.

11.3 Consequences for the original hypothesis

The original hypothesis H_{iso} remains NOT ESTABLISHED. The pivot does not refute it; it identifies it with a strong bridge assumption and observes that no evidence in either project supports that bridge. Two directions follow. If a continuity or Hölder bridge can be justified on independent grounds for a specified Π , H_{iso} would follow from the gap by Theorem 1(b),(e). If such a bridge fails, structural exclusion no longer follows from the gap alone; $H_{\text{acc}}^{\text{lim}}$ then remains an independent proposition that must be established directly, because continuity is sufficient but not necessary (Remark 1). In symbols, $\text{bridge} + \text{gap} \Rightarrow H_{\text{iso}}$, whereas $\text{bridge failure} \not\Rightarrow H_{\text{acc}}^{\text{lim}}$.

11.4 Outside Yang–Mills

The theorem is not specific to gauge theory. Any system in which a structural distance and a spectral distance are defined independently—a lattice model, a quantum graph, a family of operators on a varying domain—is subject to the same dichotomy. The separation of nearness notions may therefore be useful as a checklist in other contexts in which a claim of the form “proximity in one space implies proximity in another” is made implicitly.

12 Limitations and falsifiability

12.1 Limitations

- (L1) **Elementary theorem.** Theorem 1 is essentially continuity. It is stated and proved for completeness and to locate the missing assumption, not as a deep result.
- (L2) **Premise not instantiated.** $H_{\text{acc}}^{\text{lim}}$ is not established; neither is a representation-independent D_{struct} .
- (L3) **Profile windows, not excitations.** The Stage 2 states are windows of a vacuum-subtracted profile along the radial coordinate, not excitation sequences in a Hilbert space. The natural reading of Stage 2 is therefore about observable-space nearness (Section 10.1).

- (L4) **Class B reconstruction.** The YM-3B inputs are digitized from published figures and reconstructed into proxies; the preregistered uncertainty procedure accounts for the reconstruction but cannot account for what the published figures omit.
- (L5) **Smoothing and gradient flow.** Smeared and flowed observables suppress short-distance structure by design; the unsmoothed Battelli control reduces but does not eliminate this concern.
- (L6) **Frozen architecture.** The coexistence claim is conditional on the frozen decision rules, thresholds (such as the decisive $|\Delta\text{AICc}| \geq 6$) and the choice of five representations. A different, equally reasonable architecture might produce different verdicts.
- (L7) **Two spectral anchors.** SPEC-I/A and SPEC-I/B give two independent positive scales, but they are both lattice results; neither is a constructive-QFT bound.
- (L8) **No Clay claim.** The present work does not construct the continuum theory, does not prove a gap, and does not prove $V \in \overline{\mathcal{A}_{\text{YM}}} \setminus \{V\}$.

12.2 What would change the conclusions

The program is stated so that specific outcomes would force revision.

- A representation in which the vacuum-subtracted structural distance saturates at a positive floor, reproduced across refinements and across independent data, would falsify the *universality* in NO-UNIVERSAL-STRUCTURAL-EXCLUSION-DETECTED (not the Stage 2 results in the other representations).
- A decisive AICc preference ($\Delta\text{AICc} \leq -6$) for a positive floor in an admissible EMT dataset would convert EMT-UNCERTAINTY-LIMITED to a floor-supported verdict.
- A refinement-stable Class B result for CORR-I (the sign of S_B preserved at every level with the preregistered interval excluding zero) would promote the original diagnostic and revive the case for the original hypothesis.
- A proof that Π is continuous at V for a physically motivated $\mathcal{Y}_{\text{spec}}$, together with a gap, would imply an exclusion neighborhood by Theorem 1(b) and contradict $H_{\text{acc}}^{\text{lim}}$.
- A measured continuum scalar mass compatible with zero in either normalization would void Stage 3.

13 Conclusion

We began with a hypothesis that the Yang–Mills mass gap forces a structurally non-crossable neighborhood around the vacuum. The program built to test it separated massive from gapless correlator controls at fixed resolution but did not establish refinement-stable non-crossability on measured data: the central verdict is NOT ESTABLISHED. Testing the complementary hypothesis—that the vacuum is structurally accessible—across five representations produced no universal positive exclusion floor, and two independent lattice results retain a strictly positive scalar scale. The combined, frozen result is STRUCTURAL-SPECTRAL-COEXISTENCE-SUPPORTED, an empirical statement within the frozen architecture.

The theorem proved here explains why this is not a contradiction: a spectral gap constrains energies, not structure, and in the original program the two were linked by an assumed bridge. Continuity at the vacuum is a sufficient bridge to structural exclusion; accessibility in the limit sense ($H_{\text{acc}}^{\text{lim}}$) under a persistent gap would force that continuity to fail; and the failure would localize in either the structure-to-spectrum map or the topology of the spectral space. What

remains is the hard specialization: to construct $\Pi : \mathcal{S}_{\text{UNNS}} \rightarrow \mathcal{Y}_{\text{spec}}$ for Yang–Mills, to establish the limit statement $H_{\text{acc}}^{\text{lim}}$ rather than the empirical $H_{\text{acc}}^{\text{emp}}$, and to determine the topology that decides the question. We state this as the open problem of the program, as the open item of the four-part separation of Section 1, and make no claim about the Clay problem.

A Exhaustive verdict record

Table 10 lists every frozen verdict and stage status of both projects, with the representation tested, the question asked, and what the verdict rules in and out.

Table 10: Every frozen verdict and stage status of UNNS_YM_GAP (G) and UNNS_VACUUM_CONTINUITY (V). Verdict strings are quoted exactly as frozen.

ID	Source	Representation	Question	Verdict	Rules in / out
G1	UNNS_YM_GAP YM-0	Formalization	Frozen hypotheses, observables, null rules, refinement logic, falsification criteria	complete for the present program	In: preregistered protocol. Out: any physical claim.
G2	UNNS_YM_GAP YM-1	Gauge-invariant glueball/lattice corpus	Is there a working corpus of gauge-invariant quantities?	sufficient for current work	In: data base for YM-2/3. Out: nothing physical.
G3	UNNS_YM_GAP YM-2A / STRUC-I	Six-node vacuum-anchored spectral ladder	Does vacuum compression (gap reduced up to $\times 1024$) leave the ladder inadmissible?	vacuum-compressed branch not rejected	Out: short ladder as non-crossability evidence.
G4	UNNS_YM_GAP YM-2A / STRUC-PERC-I	Same ladder, four deformation families	Is fragmentation specific to approach toward the vacuum?	directionality test failed; representation/test pair falsified	Out: connectivity collapse $\equiv$ vacuum non-crossability.
G5	UNNS_YM_GAP YM-2B	CORR-I correlator chamber (AT20)	Does R_{contract} separate massive from gapless adversaries at fixed resolution?	VACUUM-BOUNDARY DIAGNOSTIC ESTABLISHED; REFINEMENT TEST REQUIRED (provisional)	In: fixed-resolution separation. Out: refinement stability.
G6	UNNS_YM_GAP YM-3A	Class C mass-anchored controls, 8 levels	Does the separation persist under common refinement?	MODEL-REFINEMENT PERSISTENT	In: model-family persistence. Out: measured-correlator persistence.
G7	UNNS_YM_GAP YM-3B	Class B Sakai-Sasaki reconstruction, 4 levels	Does the separation persist for measured data?	DATA-INDETERMINATE	In: central crossing at $\beta = 6.71$. Out: stable falsification; stable persistence.
G8	UNNS_YM_GAP YM-4	Bridge theorem under original premise	—	not entered / cancelled	Out: any structural-to-Hamiltonian bridge.
G9	UNNS_YM_GAP closure	Whole project	Is a refinement-stable non-crossable vacuum boundary established?	NOT ESTABLISHED	Out: universal non-crossability as established; not a proof of its negation.
V0	UNNS_VACUUM_CONTINUITY Stage 1	Ontology $V \in \overline{\mathcal{A}_{\text{YM}}} \setminus \{V\}$; VREF-I	What is meant by “vacuum belongs to the substrate”; failure criteria F1–F7	FROZEN (definitions; no empirical verdict)	Out: “basin”, “attractor”, “fixed point”.
V1	Baker 2022 [22] VREF-I/A	Common-refinement profiles, 3 resolutions	Is the refinement ordering interpolation-independent?	CALIBRATION-SENSITIVE	In: methodological warning. Out: continuity or isolation verdict.
V2	CEA 2017 [21] VREF-I/B	Connected field profile, deep tail	Zero floor vs positive floor; cross-refinement agreement	REFINEMENT-INDETERMINATE	In: vacuum-directed path, no floor detected. Out: absolute refinement-stable calibration; isolation.

ID	Source	Representation	Question	Verdict	Rules in / out
V3	CEA 2017 STRUC bridge	STRUC-I / STRUC-PERC-I on 5 ladders	Does refinement change regime or topology?	ARCHITECTURE-PRESERVED	In: structural organization stable. Out: absolute-amplitude agreement.
V4	Battelli 2019 [23]	Unsmoothed connected correlator, 3 ladders	Is the architecture an artifact of smearing?	UNSMOOTHED-ARCHITECTURE-CONSISTENT	In: corroboration. Out: independent primary distance measurement.
V5	Yanagihara 2019 [24] EMT-I, $R = 0.69$ fm	Vacuum-subtracted 4-component EMT norm	Zero floor vs positive floor (decisive $ \Delta\text{AICc} \geq 6$)	EMT-UNCERTAINTY-LIMITED	In: vacuum-directed, zero floor favored numerically. Out: decisive floor exclusion; EMT-CONTINUITY-CONSISTENT.
V6	Yanagihara 2019 geometry	Same at $R = 0.46, 0.92$ fm	Is the result geometry-dependent?	GEOMETRY-ROBUST-NO-FLOOR	In: no floor at any tested separation. Out: floor-free proof.
V7	Stage 2 synthesis	All Stage 2 branches	Is a universal refinement-stable positive floor detected?	NO-UNIVERSAL-STRUCTURAL-EXCLUSION-DETECTED	Out: universal non-crossability in tested representations.
V8	Stage 2 interpretation	Same	What does V7 support?	STRUCTURAL-VACUUM-ACCESSIBILITY-SUPPORTED	In: accessibility (empirical). Out: $V \in \overline{\mathcal{A}_{\text{YM}}}$ as a theorem.
V9	SPEC-I/A: AT20 [16]	$M_{0++} / \sqrt{\sigma} = 3.405(21)$	Is the continuum scalar scale positive?	SPECTRAL-SCALE-POSITIVE	In: $d_{\text{spec}} > 0$ in $\sqrt{\sigma}$ units. Out: any structural statement.
V10	SPEC-I/B: Chen et al. [15]	$r_0 M_{0++} = 4.16(11)(4)$	Same, independent source and scale	SPECTRAL-SCALE-POSITIVE	In: independent positivity. Out: conversion between $\sqrt{\sigma}$ and r_0 .
V11	BRIDGE-III	Frozen statuses V8, V9, V10	Do accessibility and spectral positivity coexist?	COEXISTENCE-SUPPORTED	In: coexistence. Out: identification; mass-gap or closure theorem.
V12	Stage 3 / project	Whole project	Project-level result	STRUCTURAL-SPECTRAL-COEXISTENCE-SUPPORTED	In: structural continuity can coexist with spectral discreteness (within frozen architecture). Out: Clay problem.

B Supporting data

B.1 YM-3A: Class C eight-level refinement

Table 11: YM-3A primary common-grid results (all matched orderings pass).

β	R_{massive}	$R_{p=8}$ matched	S_{fixed}	S_{match}
5.6924	0.018556	0.041575	0.708018	0.806712
5.8	0.012377	0.040156	1.112950	1.176904
5.8941	0.010214	0.039234	1.305025	1.345755
5.99	0.008840	0.038463	1.449513	1.470402
6.0625	0.008297	0.038108	1.512884	1.524494
6.235	0.007950	0.037863	1.555686	1.560841
6.338	0.007985	0.037888	1.551275	1.557102
6.5	0.007688	0.037668	1.589194	1.589194

B.2 YM-3B: Class B four-level refinement

Table 12: YM-3B central results and preregistered sensitivity.

β	R_B	$R_{p=8}$	S_B	S_B^{PCHIP}	95% interval	matched order
6.2	0.088576	0.092825	0.020350	0.019648	[0.004053, 0.035659]	pass
6.4	0.082775	0.089902	0.035867	0.036198	[0.017167, 0.053094]	pass
6.71	0.092700	0.091978	-0.003395	-0.003948	[-0.027454, 0.019320]	fail
6.93	0.082333	0.092774	0.051851	0.052060	[0.024726, 0.081194]	pass

Nested ladder lengths (L5/L6/L7) at $\beta = 6.71$: $S_B = -0.007663, -0.006666, -0.003395$; at the other three β all lengths are positive.

B.3 Baker calibration windows

Table 13: Ratio $C = R_{23}/R_{12}$ of successive refinement distances by window and interpolation.

window	$[u_0, u_1]$	C_{linear}	C_{PCHIP}	class (lin/PCHIP)
0	[0, 0.5]	0.449	0.357	<1 / <1
1	[0.125, 0.625]	0.546	0.275	<1 / <1
2	[0.25, 0.75]	0.532	0.544	<1 / <1
3	[0.375, 0.875]	0.987	1.036	<1 / >1
4	[0.5, 1]	1.485	1.529	>1 / >1

Table 14: CEA tail amplitude of the two physical refinements (R1, R2 common-linear), GeV^2 as recorded.

x (fm)	E_{R1}	E_{R2}	E_{R2}/E_{R1}
0.381	0.1242	0.1330	1.071
0.476	0.0726	0.0810	1.116
0.571	0.0409	0.0466	1.137
0.666	0.0222	0.0258	1.161
0.761	0.0117	0.0139	1.184

Table 15: Eight AT20 ensembles: A_1^{++} ground-state mass in lattice units and in $\sqrt{\sigma}$ units.

level	β	lattice	$a\sqrt{\sigma}$	aM	$M/\sqrt{\sigma}$
1	5.6924	$10^3 \times 16$	0.39990	0.9870(90)	2.468(42)
2	5.8	$12^3 \times 16$	0.31666	0.9080(120)	2.867(38)
3	5.8941	$14^3 \times 16$	0.26118	0.7991(92)	3.060(35)
4	5.99	$18^3 \times 18$	0.21982	0.7045(65)	3.205(32)
5	6.0625	$20^3 \times 20$	0.19472	0.6365(43)	3.269(24)
6	6.235	$26^3 \times 26$	0.15003	0.4969(29)	3.312(20)
7	6.338	$30^3 \times 30$	0.12928	0.4276(37)	3.308(29)
8	6.5	$38^3 \times 38$	0.10383	0.3474(22)	3.346(23)

B.4 CEA refinement mismatch

B.5 AT20 lattice ladder

C Reproducibility and provenance

All structural and spectral inputs of this paper are the frozen records of the two projects; no number was re-fitted for the paper. The analytics page built from the same records verified 21 verdict strings and audited the frozen-manifest hashes below. SHA-256 digests:

Item	SHA-256
CORR-I v0.1.2 chamber	f972c9d460cebc19ef165e66c03f971411062eba093876406ed2efe614e2deb2
BRIDGE-III input STAGE2_STATUS.md	991476088907af31c6d0256a992f79780e51eeef9e163d006e50ea4246e16582
BRIDGE-III input SPEC_I_A_RESULT.md	3d0d60dd827fdc953c684d8eaca0b8b9bd7d9ac14763d321a8bc39465972ba54
BRIDGE-III input SPEC_I_B_RESULT.md	0318aa14e92058facab9aa2527d1a4f34bc846400453f0f67ade0b5fd8be6caf
Chen06 source PDF	66973cf81063071cf5c58d93cff86d4778b7d8f11e83eb1f3b89a21be3047f22
AT20 published (SPEC-I/A)	af40f084cf4f4624224af726d0bb4822e449b161b03ce38972077023cbd0ca7c
AT20 arXiv v2 (SPEC-I/A)	f0eb6406bf49a00d376c0be25917173426c826050046a1bf0f32987895de23d0

Frozen-manifest audit: 9 of 9 Stage 2 manifest entries matched their recorded digests; BRIDGE-III's three input digests matched (Section 9). BRIDGE-III run timestamp (UTC): 2026-10-05T22:43:34.

Baker provenance closure. The three Baker et al. VREF-I/A profile inputs were verified byte-for-byte against the anc data files identified by the publisher as the data tables underlying the published plots (Eur. Phys. J. C **82**, 937, Data Availability). Their SHA-256 digests match the frozen VREF-I/A configuration exactly; row counts (33/25/33) and transverse supports (± 1.022434 , ± 0.511477 , ± 0.511238 fm) reproduce the frozen report. The 2023 erratum changes only the systematic-error notation for $\sqrt{F}$ at $\beta = 6.769$ in the force table and does not alter the profile data or geometry used here. Baker-source provenance is therefore closed.

File	SHA-256
EX_NP_xl0.128_beta6.240.dat	7188b65ff3ad25c60d7996fd96d69bab8f246eb6a960d4c442fb102b5b56d47e
EX_NP_xl0.128_beta6.544.dat	d72f075c6cc3d0dab00da9175101a53658f807cbdc9fc74563dac604747ebb59
EX_NP_xl0.128_beta6.769.dat	c572706572cc06aeb5b163306e3beeeecd37a0f27ff8c76687ad7d6dff194a39

References

- [1] A. Jaffe and E. Witten, “Quantum Yang–Mills theory,” in *The Millennium Prize Problems*, Clay Mathematics Institute (2006), pp. 129–152.
- [2] G. ’t Hooft, “On the phase transition towards permanent quark confinement,” Nucl. Phys. B **138** (1978) 1–25.
- [3] G. ’t Hooft, “Topology of the gauge condition and new confinement phases in non-Abelian gauge theories,” Nucl. Phys. B **190** (1981) 455–478.
- [4] L. Del Debbio, M. Faber, J. Greensite and Š. Olejník, “Center dominance and Z_2 vortices in SU(2) lattice gauge theory,” Phys. Rev. D **55** (1997) 2298–2306.
- [5] A. A. Belavin, A. M. Polyakov, A. S. Schwartz and Yu. S. Tyupkin, “Pseudoparticle solutions of the Yang–Mills equations,” Phys. Lett. B **59** (1975) 85–87.
- [6] T. Schäfer and E. V. Shuryak, “Instantons in QCD,” Rev. Mod. Phys. **70** (1998) 323–425.
- [7] T. C. Kraan and P. van Baal, “Periodic instantons with non-trivial holonomy,” Nucl. Phys. B **533** (1998) 627–659.
- [8] R. Alkofer and L. von Smekal, “The infrared behaviour of QCD Green’s functions: Confinement, dynamical symmetry breaking, and hadrons as relativistic bound states,” Phys. Rept. **353** (2001) 281–465, arXiv:hep-ph/0007355.
- [9] K. G. Wilson, “Confinement of quarks,” Phys. Rev. D **10** (1974) 2445–2459.
- [10] M. Creutz, “Monte Carlo study of quantized SU(2) gauge theory,” Phys. Rev. D **21** (1980) 2308–2315.
- [11] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice*, Lecture Notes in Physics **788**, Springer (2010).
- [12] M. J. Teper, “Glueball masses and other physical properties of SU(N) gauge theories in D=3+1: a review of lattice results for theorists,” arXiv:hep-th/9812187.
- [13] C. J. Morningstar and M. Peardon, “The glueball spectrum from an anisotropic lattice study,” Phys. Rev. D **60** (1999) 034509.
- [14] B. Lucini, M. Teper and U. Wenger, “Glueballs and k-strings in SU(N) gauge theories: calculations with improved operators,” JHEP **06** (2004) 012.
- [15] Y. Chen *et al.*, “Glueball spectrum and matrix elements on anisotropic lattices,” Phys. Rev. D **73** (2006) 014516, arXiv:hep-lat/0510074.

- [16] A. Athenodorou and M. Teper, “The glueball spectrum of SU(3) gauge theory in 3+1 dimensions,” JHEP **11** (2020) 172, arXiv:2007.06422.
- [17] K. Sakai and S. Sasaki, “Glueball spectroscopy in lattice QCD using gradient flow,” Phys. Rev. D **107** (2023) 034510, arXiv:2211.15176.
- [18] R. Sommer, “A new way to set the energy scale in lattice gauge theories and its application to the static force and α_s in SU(2) Yang–Mills theory,” Nucl. Phys. B **411** (1994) 839–854.
- [19] G. S. Bali, “QCD forces and heavy quark bound states,” Phys. Rept. **343** (2001) 1–136.
- [20] J. Greensite, “The confinement problem in lattice gauge theory,” Prog. Part. Nucl. Phys. **51** (2003) 1–83.
- [21] P. Cea, L. Cosmai, F. Cuteri and A. Papa, “Flux tubes in the QCD vacuum,” Phys. Rev. D **95** (2017) 114511, arXiv:1702.06437.
- [22] M. Baker, V. Chelnokov, L. Cosmai, F. Cuteri and A. Papa, “Unveiling confinement in pure gauge SU(3): flux tubes, fields, and magnetic currents,” Eur. Phys. J. C **82** (2022) 937, arXiv:2207.08797; Erratum: Eur. Phys. J. C **83** (2023) 1063.
- [23] N. Battelli and C. Bonati, “Color flux tubes in SU(3) Yang–Mills theory: An investigation with the connected correlator,” Phys. Rev. D **99** (2019) 114501, arXiv:1903.10463.
- [24] R. Yanagihara, T. Iritani, M. Kitazawa, M. Asakawa and T. Hatsuda, “Distribution of stress tensor around static quark–anti-quark from Yang–Mills gradient flow,” Phys. Lett. B **789** (2019) 210–214.
- [25] M. Lüscher, “Properties and uses of the Wilson flow in lattice QCD,” JHEP **08** (2010) 071.
- [26] H. Suzuki, “Energy-momentum tensor from the Yang–Mills gradient flow,” PTEP **2013** (2013) 083B03, arXiv:1304.0533.
- [27] M. Lüscher and U. Wolff, “How to calculate the elastic scattering matrix in two-dimensional quantum field theories by numerical simulation,” Nucl. Phys. B **339** (1990) 222–252.
- [28] K. Osterwalder and R. Schrader, “Axioms for Euclidean Green’s functions,” Commun. Math. Phys. **31** (1973) 83–112.
- [29] K. Osterwalder and R. Schrader, “Axioms for Euclidean Green’s functions II,” Commun. Math. Phys. **42** (1975) 281–305.
- [30] J. Glimm and A. Jaffe, *Quantum Physics: A Functional Integral Point of View*, 2nd ed., Springer (1987).
- [31] J. A. Shohat and J. D. Tamarkin, *The Problem of Moments*, American Mathematical Society (1943).
- [32] N. I. Akhiezer, *The Classical Moment Problem*, Oliver & Boyd (1965).
- [33] H. Akaike, “A new look at the statistical model identification,” IEEE Trans. Automat. Control **19** (1974) 716–723.
- [34] C. M. Hurvich and C.-L. Tsai, “Regression and time series model selection in small samples,” Biometrika **76** (1989) 297–307.
- [35] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference*, 2nd ed., Springer (2002).
- [36] F. N. Fritsch and R. E. Carlson, “Monotone piecewise cubic interpolation,” SIAM J. Numer. Anal. **17** (1980) 238–246.